

How do patents affect follow-on innovation?

Evidence from the human genome*

Bhaven Sampat Heidi L. Williams
Columbia and NBER MIT and NBER

July 10, 2014

Preliminary draft: Comments welcome

Abstract

In markets where innovation is cumulative — in the sense that discoveries are inputs into follow-on discoveries — optimal patent policy in part depends on how patents on existing technologies affect follow-on innovation. In this paper, we investigate whether patents on one specific technology — human genes — have affected follow-on scientific research and product development. Using administrative data on successful and unsuccessful patent applications submitted to the US Patent and Trademark Office, we link the exact gene sequences claimed in each patent application with data measuring gene-related scientific research (publications) and commercial investments (clinical development). We document that patented genes are positively selected relative to non-patented genes based on pre-application measures of scientific and commercial value, thus motivating two quasi-experimental approaches. First, we present a simple comparison of follow-on innovation across genes claimed in successful versus unsuccessful patent applications. Second, we use the “leniency” of the assigned patent examiner as an instrumental variable for whether the patent application was granted a patent. Both approaches suggest that — on average — gene patents have had no effect on follow-on innovation. Our confidence intervals on this estimate are meaningfully precise: we are able to reject the effect sizes estimated in past papers investigating the effect of non-patent forms of intellectual property on follow-on innovation.

*Nick Bloom, Tim Bresnahan, Joe Doyle, Dan Fetter, Amy Finkelstein, Alberto Galasso, Aaron Kesselheim, Pat Kline, Mark Lemley, Petra Moser, Jim Poterba, Arti Rai, Mark Schankerman, Scott Stern, and seminar participants at Duke, Northwestern Kellogg, the NBER Productivity lunch, Stanford, and UC-Santa Barbara provided very helpful comments. We are very grateful to Ernie Berndt for help with accessing the Pharmaprojects data; to Osmat Jefferson, the CAMBIA Lens initiative, Lee Fleming, and Guan-Cheng Li for sharing USPTO-related data; and to Lizi Chen, Cirrus Foroughi, Yunzhi Gao, Grant Graziani, Kelly Peterson, Lauren Russell, Sophie Sun, Nicholas Tilipman, and Hanwen Xu for excellent research assistance. Research reported in this publication was supported by the National Institute on Aging and the NIH Common Fund, Office of the NIH Director, through Grant U01-AG046708 to the National Bureau of Economic Research (NBER); the content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or NBER. Financial support from NIA Grant Number T32-AG000186 to the NBER, NSF Grant Number 1151497, and the NBER Innovation Policy and the Economy program is also gratefully acknowledged. Contact: bns3@columbia.edu, heidiw@mit.edu.

1 Introduction

Competitive markets may under-incentivize innovation due to the public good nature of new ideas. Intellectual property rights, such as patents, aim to address this under-investment problem by allowing inventors to capture a higher share of the social returns to their research investments. By awarding inventors a temporary right to exclude others from marketing their invention, patents aim to allow inventors to earn quasi-rents — temporarily — as a way to re-coup their research and development costs, thus providing dynamic incentives for investments in new technologies. Dating back at least to analyses such as Nordhaus (1969), optimal patent policy design has traditionally been framed as a trade-off between this benefit of providing incentives for the development of new technologies and the cost of deadweight loss from higher prices during the life of the patent.

Nordhaus-style models of optimal patent policy design have traditionally modeled innovations as isolated discoveries. However, in practice many — or most — innovations are cumulative, in the sense that any given discovery is also an input into later follow-on discoveries. When innovation is cumulative, optimal patent policy design also depends on how patents on existing technologies affect follow-on innovation. One example that has been prominent in recent policy debates is a patented human gene sequence. Sequenced genetic data is a research input into subsequent follow-on discoveries: by analyzing sequenced genetic data scientists may discover links between genetic variations and diseases, and such knowledge can be applied to commercialize medical technologies such as pharmaceutical treatments and diagnostic tests. If patents on discoveries such as human genes affect follow-on innovation, that effect is a key additional input into optimal patent policy design. In this paper, we investigate whether patents on human genes have affected follow-on scientific research and product development. Our broad goal is to inform whether the Nordhaus-style trade-off between *ex ante* incentives and deadweight loss is sufficient for optimal patent policy design, or whether — at least in this context — the effects of patents on follow-on innovation need to be considered.

A well-developed theoretical literature has documented ambiguous predictions for how patents will affect follow-on innovation. On one hand, if innovation takes place in multiple firms — for example, if one firm patents a human gene and a second firm discovers a gene-based pharmaceutical treatment that requires licensing that gene — then follow-on innovation may be lower than socially optimal due to the problem of how to divide the profit across firms (Scotchmer, 1991; Green and Scotchmer, 1995; Scotchmer, 2004). This profit division problem could be exacerbated if transaction costs hinder cross-firm licensing agreements (Heller and Eisenberg, 1998; Anand and Khanna, 2000; Bessen, 2004). Consistent with this view is an oft-cited anecdote that pharmaceutical firm Bristol Myers abandoned research on more than fifty cancer-related proteins due to conflicts with gene patent holders (Pollack, 2001). On the other hand, survey research by Cohen, Walsh, and colleagues has suggested that academic researchers adopt ‘working solutions’ to patents — including but not limited to patent infringement — such that

patents do not hinder follow-on innovation (Walsh, Arora and Cohen, 2003a,b; Cohen and Walsh, 2008). Others have argued that patents are also often ignored by private firms (Lemley, 2008). Finally, work dating back at least to Kitch (1977) has argued patents may facilitate investment and technology transfers across firms (Arora, 1995; Arora, Fosfuri and Gambardella, 2001; Kieff, 2001; Gans, Hsu and Stern, 2002, 2008), which may increase incentives for follow-on research and commercialization. Taken together, this motivates the need for empirical evidence to inform the question of how patents affect follow-on innovation.

The concern that patents may hinder follow-on innovation has been raised in a variety of markets — ranging from software to biomedical research tools to business methods — yet little empirical evidence exists to either support or refute that view. Two papers documented evidence that intellectual property rights discourage follow-on innovation in biomedical research contexts (Murray et al., 2008; Williams, 2013). However, both papers analyzed case studies that do not closely resemble broadly used forms of intellectual property, and survey evidence strongly suggests that patents as conventionally used will not discourage follow-on biomedical innovation, in part because patents generally preserve open access to materials at the exploratory research stage (Walsh et al., 2003a,b).¹ Despite this tension, policy makers have — understandably — looked to these available case studies for guidance, and the assumption that patents hinder follow-on innovation has been influential in shaping policy decisions. For example, Kesselheim, Cook-Deegan, Winickoff and Mello (2013) argue that the assumption that patents hinder follow-on innovation informed a series of recent US Supreme Court decisions that moved to restrict the set of discoveries eligible for patent protection.²

The lack of empirical evidence on this question reflects two challenges. First, measuring follow-on innovation is difficult, particularly for non-patented technologies. In most markets, it is difficult if not impossible to link technologies as embodied in patents with data tracing the development and use of those technologies in “real-world” markets. Second, a selection bias challenge arises because we may expect inventors to be more likely to file for and obtain patents on technologies that are more valuable. This type of selection raises the concern that any measured differences in follow-on innovation across patented and non-patented technologies may in part reflect

¹As we discuss more in Section 6, Dupont’s intellectual property rights on genetically engineered mice (analyzed in Murray et al. (2008)) and Celera’s intellectual property rights on their version of the sequenced genome (analyzed in Williams (2013)) both reduced “openness,” in the sense of inducing higher costs of accessing the initial discovery. The theoretical work of Aghion, Dewatripont and Stein (2008) and the evidence in Furman and Stern (2011) suggests openness may have important effects on follow-on innovation. Importantly — as stressed by Cohen and Walsh (2008) — patents as commonly used generally retain openness, and hence a priori may be expected to have sharply different effects from the intellectual property rights analyzed in the Dupont and Celera cases.

²Specifically, Kesselheim and co-authors highlight three decisions since 2010 in which the US Supreme Court has expanded the set of discoveries excluded from patentability. First, in *Bilski v. Kappos* the Court invalidated patent claims on an investment strategy, announcing it supported a “high enough bar” on patenting abstract ideas that it would not “put a chill on creative endeavor and dynamic change.” Second, in *Mayo v. Prometheus*, the Court invalidated patent claims on methods of using genetic variation to guide pharmaceutical dosing, expressing concern that “patent law not inhibit further discovery by improperly tying up the future of laws of nature.” Finally, in *AMP v. Myriad* (as we will discuss more), the Court invalidated patent claims on genomic DNA, concluding that “[g]roundbreaking, innovative, or even brilliant” discoveries of natural phenomena are not patentable given that such patents “would ‘tie up’ the use of such tools and thereby inhibit future innovation premised upon them.” Since the publication of Kesselheim, Cook-Deegan, Winickoff and Mello (2013), in *Alice Corp v. CLS Bank* the Court invalidated patent claims on software, arguing the claims represent abstract ideas and that implementing the claims on a computer “add[s] nothing of substance to the underlying abstract idea” and is thus also patent ineligible.

the selection of which technologies are patented.

The contribution of this paper is to construct new data and develop two new quasi-experimental approaches to investigate how patents on one type of discovery — human genes — have affected follow-on scientific research and product development. Our data construction builds on administrative records from the US Patent and Trademark Office (USPTO) on the census of patent applications, both successful and unsuccessful. We take advantage of the fact that patent applications claiming human genes as intellectual property must disclose the exact DNA sequences claimed in the text of the patent. By applying bioinformatics methods (Jensen and Murray, 2005), these DNA sequences can be annotated with gene identifiers, and these gene identifiers can in turn be linked to standard scientific and medical datasets providing measures of follow-on scientific research and product development related to each gene. Specifically, we measure the scientific publications related to each gene as an indicator of scientific research investments, and measure the use of genes in pharmaceutical clinical trials and diagnostic tests as indicators of commercial research investments. The key advantage of this DNA sequence-based match is that it enables the construction of a unique dataset which precisely links the technologies embodied in patent applications (namely, DNA sequences) with “real-world” measures of the same technologies.

These data allow us to document new facts relating patents, scientific research, and commercial investments on the human genome. We document clear evidence of selection into patenting: genes claimed in granted patent applications appear more valuable both scientifically (as measured by publications) and commercially (as measured by clinical trials) prior to receiving patents, relative to genes that are never claimed in granted patent applications.

Motivated by this evidence of selection, we develop two quasi-experimental methods for estimating how patents have affected follow-on innovation. First, we present a simple comparison of follow-on innovation across genes claimed in successful versus unsuccessful patent applications. This analysis is in the spirit of Bound’s (1989) analysis of the labor supply effects of disability insurance, which used disability insurance applicants who failed the medical screening as a control group for beneficiaries. In our context, this method is reasonable as long as — conditional on being included in a patent application — whether a patent application is granted a patent is as good as random. Empirically, genes claimed in successful and unsuccessful patent applications appear similar on characteristics fixed at the time of application, providing some evidence for the validity of this empirical approach.

Second, we develop a novel instrumental variable for which patent applications are granted patents: namely, the “leniency” of the assigned patent examiner. Patent examiners are charged with a uniform mandate: grant patents to patent-eligible, novel, non-obvious, and useful inventions. However, prior research has documented that in practice this mandate leaves patent examiners a fair amount of discretion (Cockburn, Kortum and Stern, 2003; Lemley and Sampat, 2010, 2012). We leverage the interaction of this across-examiner heterogeneity with the quasi-random assignment of patent applications to examiners as a source of variation in which patent applications are granted patents. Past qualitative evidence — and new empirical evidence we document — supports the assertion

that the assignment of patent applications to examiners is plausibly random conditional on some covariates (such as application year and technology type), suggesting that the leniency of the patent examiner to which a patent application is assigned can provide a valid instrument for whether the patent application is granted a patent.

In contrast with what one would infer from a naive comparison of follow-on innovation on patented and non-patented genes, both of our quasi-experimental approaches suggest that gene patents have not had quantitatively important effects on either follow-on scientific research or follow-on commercial investments. The estimates from our first quasi-experimental approach — comparing follow-on innovation across genes claimed in successful versus unsuccessful patent applications — document economically small and meaningfully precise estimates. For our publications outcome, our 95% confidence interval can reject declines greater than 7% and reject increases of more than 0.2%. For our clinical trials outcomes, our 95% confidence interval can reject declines greater than 3.5% and reject increases of more than 2.7%. For our diagnostic test outcome, our 95% confidence interval can reject declines greater than 17% and reject increases of more than 1.5%.

Given that the previous literature (Murray et al., 2008; Williams, 2013) documented evidence that intellectual property rights reduced follow-on innovation outcomes on the order of 30%, our estimates suggest meaningfully smaller effects in this context. Our estimates are also consistent with the recent contemporaneous contribution of Galasso and Schankerman (2013), which investigates how patents affect follow-on innovation using variation in patent invalidations generated by the quasi-random assignment of court cases to judges at the US Court of Appeals for the Federal Circuit. In their analysis of heterogeneous treatment effects across technology types, they find no evidence that patent invalidations increase patent citations - or clinical trials - in the technology field which includes most of our gene patents.³

Our first quasi-experimental approach has the advantage of generating meaningfully precise estimates, but has the cost of relying on a less transparent source of variation. Our second quasi-experimental approach — using the leniency of the assigned patent examiner as an instrument for which patent applications are granted patents — has the advantage of relying on a clean and transparent source of variation, but generates less precise estimates (although even these estimates can usually reject the 30% declines documented in the prior literature). Taken together, the fact that these two approaches generate similar conclusions gives us confidence in focusing our interpretation on the more precise estimates generated by the first approach.

Our empirical estimates speak to a recent set of high-profile legal rulings on the case *Association for Molecular Pathology v. Myriad Genetics*. The firm Myriad Genetics was granted patents on human genes correlated with risks

³Specifically, the most common patent class in our sample is 435 (Chemistry: molecular biology and microbiology), which is included in the “Drugs” subcategory of the NBER patent classification (<https://sites.google.com/site/patentdataprotect/Home/downloads/patn-data-description>); Galasso and Schankerman do not find evidence that patent invalidations increase follow-on innovation in that subcategory. Unfortunately, their variation is not directly applicable in our context because very few human gene patents have ever been asserted in court (Holman, 2007, 2008).

of breast and ovarian cancer, and in June 2013 the US Supreme Court unanimously ruled to invalidate a subset of Myriad’s gene patent claims. Importantly for this paper, although gene patents were controversial for a variety of reasons — including concerns about the ethics of “owning” human genes — in the end the Court decision argued that DNA patents “*would ‘tie up’...such tools and...inhibit future innovation.*” As argued by a 2006 National Academies report and by Caulfield et al. (2013), prior to this paper there was little empirical evidence available to either support or refute the assertion that gene patents were impeding follow-on innovation.⁴

Taken together, these results suggest two conclusions. First, for the case of human genes, the traditional patent trade-off of ex ante incentives versus deadweight loss may be sufficient to analyze optimal patent policy design, as the effects of patents on follow-on innovation appear to be quantitatively small. Second, the available evidence suggests there may be important heterogeneity in how and when intellectual property policies have effects on follow-on innovation.⁵ A variety of evidence suggests that policies governing access to materials appear to have important effects on follow-on innovation (Murray et al., 2008; Furman and Stern, 2011; Williams, 2013), and the evidence from Galasso and Schankerman (2013) suggests that patents may deter follow-on innovation in some specific contexts. This growing body of empirical evidence suggests policy-relevant distinctions that should inform future theoretical and empirical research on cumulative innovation.

Methodologically, our two quasi-experimental approaches build on similar applications in labor economics and public finance. Our comparison of accepted and rejected patent applications builds on Bound (1989) and von Wachter, Song and Manchester (2011)’s investigations of the disincentive effects of disability insurance on labor supply, and Aizer, Eli, Ferrie and Lleras-Muney (2013)’s investigation of the effects of cash transfers on mortality. Likewise, our approach of using examiner “leniency” as a source of variation in which patent applications are granted patents builds on past work investigating the effects of incarceration length using variation across judges (Kling, 2006), the effects of foster care using variation across foster care case workers (Doyle, 2007, 2008), and the disincentive effects of disability insurance on labor supply using variation across disability insurance examiners (Maestas, Mullen and Strand, 2013). Especially given the relatively small number of patent law changes in countries like the US in recent years, these two sources of quasi-experimental variation may provide valuable opportunities to investigate the effects of patents in a variety of other applications.

Section 2 describes our data. Section 3 documents evidence of selection in which genes are patented, motivating our two quasi-experimental approaches. Section 4 presents estimates from our first quasi-experimental ap-

⁴Despite not directly investigating the effects of gene patents, the Williams (2013) evidence — documenting that a non-patent form of database protection on the human genome hindered follow-on innovation — was cited in several briefs filed in the *AMP v. Myriad* case, such as http://www.americanbar.org/content/dam/aba/publications/supreme_court_preview/briefs-v2/12-398_pet_authcheckdam.pdf and <http://www.yaleisp.org/wp-content/uploads/2012/01/31253-BRCA-AMICUS-BRIEF-FINAL.pdf>; see Marshall (2013) for a recent discussion.

⁵We here do not focus on the estimates from Murray and Stern (2007), who documented evidence that patents hindered subsequent citations to matched scientific papers (a measure of follow-on innovation), because subsequent work by the authors (Fehder, Murray and Stern, 2012) documented that this result was specific to the first few years after the founding of the journal they analyzed.

proach, comparing follow-on innovation across genes claimed in successful and unsuccessful patent applications. Section 5 presents estimates from our second quasi-experimental approach, using the leniency of the assigned patent examiner as an instrumental variable for whether the patent application was granted a patent. Section 6 interprets our empirical results in light of the growing literature investigating the impact of intellectual property rights in markets where innovation is cumulative, and Section 7 concludes.

2 Data

This section describes our data construction.⁶ To fix ideas, we start by describing an example patent application — USPTO patent application 08/483,554 — claiming intellectual property over the BRCA1 gene, and describe our measures of follow-on innovation in the context of that example (Section 2.1). We then describe in more detail how we construct our sample of USPTO patent applications claiming intellectual property over human genes (Section 2.2), and our gene-level measures of follow-on scientific research and product development (Section 2.3).

2.1 Example: USPTO patent application 08/483,554

On 7 June 1995, Myriad Genetics filed USPTO patent application number 08/483,554: *17Q-Linked Breast and Ovarian Cancer Susceptibility Gene*. This application was subsequently granted a patent on 5 May 1998 (US patent number 5,747,282), and would later become a focus of the US Supreme Court case *Association for Molecular Pathology v. Myriad Genetics*.

This patent claimed intellectual property over an isolated sequence of nucleotide bases (adenine, cytosine, guanine, and thymine), the sequence of which is listed explicitly in the patent (*SEQ ID NO:1: ATG GAT TTA...*). By comparing this sequence of nucleotide bases to the census of human gene sequences, we can uncover that this sequence corresponds to the BRCA1 gene.

The text of the Myriad patent describes how variation in the precise sequence of nucleotide bases in the BRCA1 gene can induce variation in an individuals' risk of developing breast and ovarian cancers. For example, women with abnormal BRCA1 or BRCA2 genes have 40 to 80 percent lifetime risk of developing breast cancer, relative to about 12 percent in the general population.⁷

Such links between genetic variation and diseases are referred to as genotype-phenotype links. In the case of the BRCA1 gene, scientific papers have investigated links between BRCA1 and breast cancer (141 publications), ovarian cancer (96 publications), and pancreatic cancer (2 publications).⁸ Once the scientific knowledge of a given genotype-phenotype link has been documented, this knowledge can be applied to develop medical technologies

⁶Appendix A describes our data construction in more detail.

⁷These statistics are drawn from GeneReviews, published by the US National Institutes of Health (NIH).

⁸These data are drawn from the NIH Online Mendelian Inheritance in Man (OMIM) database, described below.

such as pharmaceutical treatments and gene-based diagnostic tests. In the case of the BRCA1 gene, more than 200 pharmaceutical clinical trials have been conducted that focus on mutations in the BRCA1 gene, and a BRCA1 genetic test is marketed by the firm Myriad Genetics.⁹ The goal of our data construction is to trace these types of measures of follow-on scientific research and product development for each human gene, and to link these data to records of which human genes have been included in USPTO patent applications and granted patents.

2.2 Constructing USPTO patent application sample

Our quasi-experimental approaches require constructing data on the census of published USPTO patent applications that claim intellectual property over human genes — both successful (granted patents) and unsuccessful (not granted patents) patent applications. Traditionally, unsuccessful USPTO patent applications were not published. However, as part of the American Inventors Protection Act of 1999, the vast majority of USPTO patent applications filed on or after 29 November 2000 are published in the public record — regardless of whether they are granted patents — eighteen months after the filing date.¹⁰

From the census of USPTO patent applications, we identify the subset of applications claiming intellectual property over human genes. To do this, we follow the methodology proposed by Jensen and Murray (2005).¹¹ In short, since the early 1990s DNA sequences have been listed in USPTO patent applications in a standard format, labeled with the text “SEQ ID NO” (sequence identification number). This standardized format allows for DNA sequences to be cleanly extracted from the full text of USPTO published patent applications. Once extracted, standard bioinformatics methods can be used to compare these sequences against the census of human gene DNA sequences in order to annotate each sequence with standard gene identifiers that can be linked to outside databases.¹²

We apply this Jensen and Murray (2005) methodology to construct two samples. First, we construct a “full sample” that aims to include the census of published USPTO patent applications that claim any DNA sequences — not just human gene sequences — in their patent claims (Bacon et al., 2006). Second, we construct a “human gene sample” that includes only the subset of published USPTO patent applications that claim human genes (as

⁹These data are drawn from the Citeline Pharmaprojects database and the NIH GeneTests.org database, described below.

¹⁰For more details, see <http://www.uspto.gov/web/offices/pac/mpep/s1120.html> and the discussion in Lemley and Sampat (2010). Most applications not published eighteen months after filing are instead published sixty months after filing.

¹¹While we focus on the Jensen and Murray (2005) definition of gene patents, there are many different types of DNA-related patent claims — e.g., protein-encoding sequences, expressed sequence tags (ESTs), single-nucleotide polymorphisms (SNPs), sequence-based claims, and method claims that pertain to specific genes or sequences; see Scherer (2002) and Holman (2012) for more discussion. As one point of comparison, nearly all of the patent applications in our sample are included in the *DNA Patent Database* (<http://dnapatents.georgetown.edu/SearchAlgorithm-Delphion-20030512.htm>), which is constructed using a completely different methodology; we are grateful to Robert Cook-Deegan and Mark Hakkarinen for sharing the patent application numbers included in the *DNA Patent Database*, which enabled this comparison. Our empirical strategies could of course be applied to any empirically implementable definition of gene patents, but we are not aware of data on any other alternative definitions.

¹²Following Jensen and Murray (2005), we focus attention on the subset of DNA sequences that are explicitly listed in the claims of the patent applications (excluding DNA sequences that are referenced in the text of the patent application — and hence listed in the SEQ ID NO format — but not explicitly claimed as intellectual property).

opposed to non-human genes, such as mouse genes; Lee et al. (2007)).¹³ This human gene sample is the focus of our analysis, given the focus in this paper on investigating how patents on human genes have affected follow-on innovation. However, this sample includes far fewer patent applications than the full sample: we have around 1,500 applications in the human gene patent sample, relative to around 30,000 applications in the full sample. This matters for our second quasi-experimental design, because it is much more difficult to reliably estimate how patent grant propensities vary across examiners in the human gene patent sample alone. Reflecting this challenge, we use the full sample to estimate the examiner patent grant propensities, and apply those estimates to the examiners in our human gene sample.

Table 1 presents patent application-level summary statistics on our full sample and our human gene sample. Our full sample includes 30,243 USPTO patent applications with DNA sequences listed in their claims, while our human gene sample includes 1,582 USPTO patent applications with human genes listed in their claims. Our data primarily includes US patent applications filed from 2000-2010, given that the requirement for unsuccessful patent applications to be published came into force for patent applications filed on or after 29 November 2000. A small share (<0.01%) of our sample has an application year prior to 2000; we retain these observations in our sample in the absence of a clear reason to exclude them. The human gene sample ends with application year 2005, marking the last year of data used as an input into the data construction done by Lee et al. (2007). In both samples, about 30 percent of the patent applications had been granted patents by 2010, although this raw statistic does not account for censoring (our estimation will account for censoring).

While our quasi-experimental approaches can only be implemented over the time period when unsuccessful patent applications were published (that is, applications filed after 29 November 2000), for the purposes of our descriptive analyses we also apply the Jensen and Murray (2005) methodology to identify granted USPTO patents that claim human genes but were filed prior to 29 November 2000.¹⁴

2.3 Measuring follow-on innovation

We collect data on three measures of gene-level follow-on innovation: scientific publications as a measure of scientific research effort; and two measures of product commercialization: gene-related pharmaceutical research, and gene-based diagnostic tests.

We collect data on the scientific publications related to each gene from the Online Mendelian Inheritance in Man (OMIM) database, which catalogs scientific papers that have documented evidence for links between genetic variation and phenotypes. Using this data, we construct a count of the number of scientific papers published related

¹³Lee et al. (2007) applied the Jensen-Murray methodology to USPTO published patent applications available through 2005 (Jensen and Murray (2005) applied their methodology only to granted USPTO patents).

¹⁴These data are also drawn from Lee et al. (2007), and include USPTO patents granted through 2005.

to each gene — across all phenotypes — in each year.

Because of the long time lags between basic drug discovery and the marketing of new drugs, new approvals of drugs that take advantage of sequenced genetic data are just barely starting to enter the market (Wade, 2010). Given these time lags, rather than using drug approvals as a measure of pharmaceutical research, we instead focus on an intermediate measure of drug discovery — namely, drug compounds currently under development, as disclosed in clinical trials. Specifically, we measure gene-related pharmaceutical clinical trials using the Citeline Pharmaprojects database, a privately-managed competitive intelligence database that tracks drug compounds in clinical trials and — critically for our project — assigns gene identifiers to compounds related to genetic variation on specific genes. Using this data, we construct a count of the number of clinical trials related to each gene in clinical trials in each year.

Finally, we collect data on gene-based diagnostic tests from the GeneTests.org database, a genetic testing registry. Some gene-based tests provide individuals with information about disease risk, such as the BRCA tests related to risks of breast and ovarian cancers; other gene-based tests identify individuals as relatively more or less appropriate for a given medical treatment, such as a test which predicts heterogeneity in the side effects of the widely-prescribed blood thinner warfarin. Using the GeneTests.org data, we construct an indicator for whether a gene is included in any gene-based diagnostic test as of 2012. Unfortunately, this data is only available in a cross-section (not a panel) and thus is less useful for our empirical analysis.

A priori, the impact of patents on follow-on scientific research could differ from the impact of patents on product development. For example, many have argued that most patented inventions are made available to academic researchers on sufficiently favorable licensing terms that academics are able to continue their research (USPTO, 2001), so that even if transaction costs hinder licensing agreements for commercial applications we could expect to see no impact of patents on measures of follow-on academic research such as scientific publications. Alternatively, patents may change researchers' incentives of whether to disclose the results of their research through academic publications (Moon, 2011), in which case observed differences in scientific publications could be explained by differences in disclosure rather than by differences in the true amount of underlying scientific research. By contrast, we would not expect the product development outcomes we measure to be affected by such disclosure preferences given that the measures we observe are revealed in the natural course of firms commercializing and selling their technologies.

3 Comparison of patented and non-patented human genes

Our interest in this paper is in comparing follow-on innovation across patented and non-patented inventions. In order to investigate this question, we need to first examine the selection process, in the sense of understanding

which inventions are patented. On one hand, we may expect inventors to be more likely to both file for and be granted patents on technologies that are more valuable: inventors may be more willing to pay the time and monetary cost of filing patent applications for inventions that are more valuable, and patent applications claiming intellectual property over more valuable inventions may also be more likely to “clear the bar” and be granted patents. On the other hand, the USPTO grants patents based on criteria — patent-eligibility, novelty, non-obviousness, and usefulness — that may not closely correspond with measures of scientific and commercial value (Merges, 1988). Hence, it is an empirical question whether inventions with higher levels of scientific and commercial value are more likely to be patented.

There is a remarkable absence of empirical evidence on this question, largely due to a measurement challenge: in most markets, it is very difficult to measure the census of inventions, and to link those inventions to patent records in order to identify which inventions are and are not patented. The only other paper we are aware of which has undertaken such an exercise is Moser (2012), who constructs a dataset of innovations exhibited at world’s fairs between 1851 and 1915 and documents that “high quality” (award-winning) exhibits were more likely to be patented, relative to exhibits not receiving awards. Given this dearth of previous estimates, documenting evidence on the selection of inventions into patenting is itself of interest. In addition, if we do observe evidence of selection into patenting, that would imply that any measured differences in follow-on innovation across patented and non-patented genes may in part reflect the selection of which genes are included in patent applications and granted patents (as opposed to an effect of patents on follow-on innovation).

For this exercise, we start with the full sample of human genes ($N=26,440$). As measured in our data, approximately 29 percent of human genes have sequences that are explicitly claimed in granted US patents ($N=7,717$).¹⁵ Figure 1 documents trends in follow-on innovation by year separately for genes that ever receive a patent (triangle-denoted solid blue series), and for genes that never receive a patent (circle-denoted dashed red series). For scientific publications (Figure 1(a)), we plot the average log number of scientific publications by year — plus one — in each year from 1970 to 2012.¹⁶ For clinical trials (Figure 1(b)), we plot the average log number of clinical trials by year — plus one — in each year from 1995 to 2011.¹⁷

Because by construction our human gene patents are measured starting in the mid-1990s (when the SEQ ID

¹⁵As a point of comparison, Jensen and Murray (2005) document that as of 2005, approximately 20 percent of human genes had sequences that were explicitly claimed as granted patents. Because our sample includes patent applications that were granted patents after 2005, we would expect our estimate to be mechanically larger.

¹⁶We focus on the average log number of scientific publications by year because even within a calendar year, the number of publications per human gene is quite right-skewed. For example, in 1970 the median number of publications per gene is 0, the mean is 0.03, and the maximum is 15. The pattern of selection that we document is unchanged if we instead plot the average number of scientific publications by year, or the share of genes with at least one scientific publication by year.

¹⁷These years — 1995 to 2011 — are the only years for which the Pharmaprojects data are available. As with the scientific publications measure, we focus on the average log number of clinical trials by year because this variable is quite right-skewed. For example, in 1995 the median number of clinical trials per gene is 0, the mean is 0.16, and the maximum is 127. Again, the pattern of selection that we document is unchanged if we instead plot the average number of clinical trials by year, or the share of genes with at least one clinical trial by year.

NO notation was introduced), the cleanest test of selection into patenting is a comparison of pre-1990 follow-on innovation across patented and non-patented genes. While the clinical trials data only starts later — in 1995 — the scientific publications data is available prior to 1990 so can be used for this comparison. Looking at the data series in Figure 1(a) from 1970 to 1990 provides clear evidence of positive selection: genes that will later receive patents were more scientifically valuable — based on this publications measure — prior to being patented. Moreover, even within the 1970 to 1990 period this positive selection appears not just in the levels of follow-on innovation, but also in the trends: genes that will later receive patents appear to have divergent trends in scientific publications relative to genes that will never receive patents, even before any patents are granted.

These patterns — a level difference, and a divergence in trends — also appear in the clinical trials data presented in Figure 1(b), although it is not possible to cleanly separate selection and treatment because that data series start in 1995. Likewise, if we tabulate the probability that genes are used in diagnostic tests for patented and non-patented genes, 13.45% of patented genes are used in diagnostic tests, compared to 6.39% of non-patented genes.

While we only observe our diagnostic test outcome in a cross-section, for our clinical trial outcomes we can replicate a version of this analysis which focuses on the set of gene patents whose applications were filed in or after 2001, in which case 1995 to 2000 can serve as a “pre-period” that isolates a selection effect. Appendix Figure D.1 documents that we see very similar patterns of selection on levels as well as pre-trends using this variation.

Taken together, these patterns suggest strong evidence of positive selection: genes that will later receive patents appear more scientifically and commercially valuable prior to being granted patents, relative to genes that will never receive patents. This evidence is important for two reasons. First, this evidence suggests that our measures of value (pre-patent filing scientific publications and clinical trials) are correlated with patenting activity, which provides some evidence that these measures can provide a meaningful basis for assessing selection in our two quasi-experimental approaches. Second, this evidence implies that a simple comparison of follow-on innovation across patented and non-patented genes is unlikely to isolate an unbiased estimate of how gene patents affect follow-on innovation. While a naive comparison that did not account for selection would conclude based on Figure 1 that patents encourage follow-on innovation, as we will see in Sections 4 and 5 our two quasi-experimental approaches will suggest a different conclusion.

4 Comparison of accepted and rejected patent applications

Our first quasi-experimental source of variation investigates a simple idea, which is whether genes that were included in unsuccessful patent applications can serve as a valid comparison group for genes that were granted patents.

A comparison of follow-on innovation on genes included in accepted and rejected patent applications will be

valid if — conditional on being included in a patent application — whether an application being granted a patent is as good as randomly assigned.¹⁸ A priori, we might expect this assumption to be invalid. The USPTO is responsible for determining whether inventions claimed in patent applications should be granted patents based on the criteria of patent-eligibility, novelty, non-obviousness, and usefulness — criteria that we would not expect to generate random assignment of patents across applications. However, as we will uncover through our analysis in Section 5, variation in the ‘leniency’ of different patent examiners appears to largely explain which gene patent applications are granted patents, suggesting that patent grants being “as good as randomly assigned” across gene patent applications may be reasonable. Empirically, we will document in this section that genes claimed in successful and unsuccessful patent applications appear similar on observable characteristics fixed at the time of application, providing some evidence for the validity of this empirical approach. While this comparison is quite simple, we will see that the resulting estimates are quite similar to the estimates from our second quasi-experimental source of variation (the examiner leniency variation, presented in Section 5).

We first provide a brief overview of the USPTO patent examination process, as background for our analysis of accepted and rejected applications (Section 4.1), and then document our estimates (Section 4.2 and 4.3).

4.1 Overview of the USPTO patent examination process

The USPTO is responsible for assessing whether patent applications should be granted patents based on four criteria: patent-eligibility, novelty, non-obviousness, and usefulness.¹⁹ To give a sense of scale, in 2010 the USPTO was staffed by over 6,000 patent examiners, and received over 490,000 applications.²⁰

A central USPTO office assigns application numbers to incoming applications, as well as patent class and subclass codes detailing the type of technology embodied in the application.²¹ These class and subclass numbers determine which of ~300 so-called Art Units — specialized groups of examiners — will review the application.²² Within an Art Unit, a supervisory patent examiner assigns the application to a patent examiner for review. We defer a discussion of this within-Art Unit assignment of applications to patent examiners to Section 5.1, because those details are most relevant to our second source of quasi-experimental variation.

While the patent application process up to this point is quite structured, from this point forward substantial discretion is left in the hands of individual examiners, who are responsible for determining which — if any —

¹⁸Strictly speaking, patent applications are never formally rejected by the USPTO, only abandoned by applicants (Lemley and Sampat, 2008). We refer to such applications colloquially as “rejected” throughout the text.

¹⁹The discussion in this section draws heavily on Cockburn, Kortum and Stern (2003), Lemley and Sampat (2012), and GAO (2005). See Appendix B for more detail on the USPTO patent examination process.

²⁰See GAO (2010) and http://www.uspto.gov/web/offices/ac/ido/oeip/taf/appl_yr.htm.

²¹There are currently over 450 patent classes, and more than 150,000 subclasses; see <http://www.uspto.gov/patents/resources/classification/overview.pdf>.

²²See <http://www.uspto.gov/patents/resources/classification/art/index.jsp>. For the current version of the class/subclass-to-Art Unit concordance, see <http://www.uspto.gov/patents/resources/classification/caau.pdf>. The main Art Units in our sample are from the 1600 group (Biotechnology and Organic Chemistry).

claims in the application are patentable. The examiner returns an initial decision on the application, which can be an allowance or — more commonly — a “non-final rejection.” For non-final rejections, applicants have a fixed length of time (usually six months) to revise the application. The examiner can then allow the application, issue a second rejection, or pursue subsequent rounds of negotiation.

4.2 Graphical analysis

For this exercise, we start with the sample of genes included in at least one patent application in our USPTO patent applications sample (N=15,530; 59% of the full sample of 26,440 human genes). Of this sample, 4,858 genes are claimed in a patent application that is subsequently granted a patent (31%), relative to 10,672 genes that are never observed to be subsequently granted a patent (69%).²³ Figure 2(a) illustrates the time pattern of when the patented group receives its (first) patent grant over time. Over half of these genes have received a patent by 2005, and (by construction) all have received a patent by 2010.

Figures 2(b) and 2(c) document trends in follow-on innovation by year. As in Figure 1, we plot the average log number of scientific publications by year — plus one — in each year from 1970 to 2012 (Figure 2(b)), and the average log number of clinical trials by year — plus one — in each year from 1995 to 2011 (Figure 2(c)). The solid blue triangle-denoted line represents genes claimed in at least one granted patent, and the dashed red circle-denoted line represents genes claimed in at least one patent application but never in a granted patent.

As a point of comparison, the dashed green square-denoted line represents genes never claimed in a patent application (N=10,910; 41% of the full sample of 26,440 human genes). Comparing this group of genes to the other two groups of genes, we see clear evidence of selection into patent filing: genes included in successful and unsuccessful patent applications are much more valuable both scientifically (publications) and commercially (clinical trials) prior to the patent application filing compared to genes that are never claimed in a patent application. In terms of interpreting the results in Section 3, this suggests that some of the selection in which genes are patented is driven by selection into which genes are included in patent applications (as opposed to being solely driven by which genes are granted patents, conditional on being included in patent applications).

The key comparison of interest in this section is across the first two data series: genes included in successful patent applications, and genes included in unsuccessful patent applications. Strikingly, we see little evidence of selection in pre-2001 levels or trends of our two follow-on innovation measures once we limit the sample to genes included in patent applications. For scientific publications (Figure 2(b)), the two groups follow each other quite closely from 1970 to 1990, and diverge slightly in trends from 1990 to 2000 — with genes that will subsequently be included in unsuccessful patent applications having slightly more scientific publications. For clinical trials (Figure

²³Note that this 4,858 figure is mechanically lower than the 7,717 figure in Section 3, because we here focus only on patents granted on patent applications filed after 29 November 2000 (the date when unsuccessful applications began to be published).

2(c)), the two groups follow each other quite closely in both levels and trends over all available pre-2001 years of data. Taken at face value, the similarity of these two groups in pre-2001 outcomes provides some evidence for the validity of this empirical approach. A priori, one might have expected genes that were more scientifically or commercially valuable to have been more likely to receive patents. However, conditional on being included in a patent application, this appears not to be the case. As we will see in Section 5, this may be because conditional on being included in a patent application, the leniency of the patent examiner to which the application is assigned appears to explain a large share of the variation in which gene patent applications are granted patents.

Looking at the post-2001 time period, we see that although these two groups of genes diverge (by construction) in whether they are claimed in granted patents (Figure 2(a)), we do not see any evidence of a divergence in follow-on innovation outcomes. That is, these figures suggest that gene patents have not had a quantitatively important effect on either follow-on scientific research or on follow-on commercialization.

4.3 Regression analysis

Table 2 quantifies the magnitudes of these differences, and presents standard errors that allow for statistical inference. Panel A of Table 2 pools all post-2001 years together; Panel B of Table 2 documents separate estimates by each individual year. For our scientific publication and clinical trial outcomes, we present OLS estimates with a log (plus one) outcome; for our diagnostic test outcome we present OLS estimates.

Interpreting the publications coefficient in Column 1 of Panel A of Table 2, genes receiving patents have a marginally statistically significant but economically small 3.4 percent fewer publications than genes included in patent applications but not granted patents over the time period 2001-2012 (p-value of 0.063). Our 95% confidence interval can reject declines of greater than 7%, and reject increases of more than 0.19%.

Interpreting the clinical trials coefficient in Column 2 of Panel A of Table 2, genes receiving patents had a statistically insignificant 0.36% percent fewer clinical trials than genes included in patent applications but not granted patents over the time period 2001-2011. Our 95% confidence interval can reject declines of greater than 3.4%, and reject increases of more than 2.7%.

While we only observe our diagnostic test outcome in a 2012 cross-section, that outcome is also very similar across the two groups: the point estimate in Column 3 of Panel A of Table 2 suggests that genes receiving patents had a statistically insignificant 0.9 percentage point decrease in the likelihood of being included in a diagnostic test as of 2012 relative to genes included in patent applications but not granted patents. Our 95% confidence interval can reject declines of greater than 2 percentage points and reject increases of more than 0.17 percentage points. Relative to a mean of 12%, this confidence interval suggests that we can reject declines in this outcome of greater

than 17%.²⁴

5 Analyzing examiner-level variation in patent grant propensities

Our second source of quasi-experimental variation aims to construct an instrumental variable for whether a patent application is granted a patent. Our key idea is to build on previous research which has established that although patent examiners are charged with a uniform mandate, in practice examiners have a fair amount of discretion, and this discretion appears to translate into substantial variation in the decisions different examiners make on otherwise similar patent applications (Cockburn, Kortum and Stern, 2003; Lichtman, 2004; Lemley and Sampat, 2010, 2012).²⁵

Of course, the existence of variation in leniency across examiners on its own is not sufficient for examiner-level leniency to be a valid instrumental variable for patent grants. First, there must exist a sufficiently strong empirical relationship between the instrument and an individual patent application's probability of being granted. We will document empirical support for this first assumption in our data. Second, the monotonicity condition requires that any patent application granted by a strict examiner would also be granted by a lenient examiner, and any patent application rejected by a lenient examiner would also be rejected by a strict examiner. While inherently untestable, we will document some empirical evidence suggesting that the monotonicity assumption is reasonable in our context. Finally, the instrument needs to satisfy the exclusion restriction: the instrument can only affect the outcome (follow-on innovation) through the likelihood that a gene is patented. As we will discuss, the institutional context suggests that the assignment of patent applications to USPTO patent examiners should be effectively random conditional on some covariates (such as application year and technology type); consistent with this assumption, we will document that genes assigned to 'lenient' and 'strict' examiners look similar on observable characteristics fixed at the time of patent application.²⁶

In this section, we provide some qualitative background on the key institutional features underlying our empirical strategy (Section 5.1), describe the construction of our measure of examiner leniency (Section 5.2), and then present our estimates (Sections 5.3, 5.4, 5.5, 5.6, and 5.7).

²⁴As a point of comparison, only 3% of genes never included in a patent application are included in a diagnostic test as of 2012.

²⁵One of the individuals interviewed by Cockburn, Kortum and Stern (2003) described this variation informally by saying, "there may be as many patent offices as there are patent examiners." Similarly, a trade publication written by a former USPTO patent examiner and current patent agent (Wolinsky, 2002) described this variation by saying: "*The successful prosecution of a patent application at the USPTO requires not only a novel invention and adequate prosecution skills, but a bit of luck...If you knew the allowance rate of your examiner, you could probably estimate your odds of getting your patent application allowed.*"

²⁶While conditional random assignment of applications to examiners assuages many potential concerns about this exclusion restriction, some additional issues remain. In particular, while we focus on variation in patent grant propensity, examiner heterogeneity may also manifest itself in other ways, such as the breadth of patent grants (in terms of the number or strength of allowed claims) and time lags in grant decisions (Cockburn, Kortum and Stern, 2003).

5.1 Assignment of patent applications to patent examiners

As discussed in Section 4.1, within an Art Unit a supervisory patent examiner assigns the application to a patent examiner for review. Because no “standard” method for the within-Art Unit assignment of applications to examiners is uniformly applied in all Art Units, Lemley and Sampat (2012) conducted written interviews with roughly two dozen current and former USPTO examiners to inquire about the assignment process. While the results of these interviews suggested that there is not a single “standard” assignment procedure that is uniformly applied in all Art Units, these interviews revealed no evidence of deliberate selection or assignment of applications to examiners on the basis of characteristics of applications other than those observed in standard USPTO datasets. For example, in some Art Units supervisors reported assigning applications to examiners based on the last digit of the application number; because application numbers are assigned sequentially in the central USPTO office, this assignment system — while not purposefully random — would be functionally equivalent to random assignment for the purposes of this study. In other Art Units, supervisors reported placing the applications on master docket based on patent classes and subclasses, with examiners specializing in those classes (or subclasses) being automatically assigned the oldest application from the relevant pool when requesting a new application. Our key conclusion from this institutional context is that effective conditional random assignment of applications to examiners — within Art Unit and application year — is plausible. Consistent with this assumption, we will document that genes assigned to ‘lenient’ and ‘strict’ examiners look similar on observable characteristics fixed at the time of patent application.

From a practical perspective, it is worth noting that informational barriers limit the extent to which we would expect patent applications to be systematically sorted across examiners in a way that would be problematic for our empirical specifications. Because of the limited attention given to patents prior to their assignment to a specific examiner, and the judgment required to determine the patentability of a given invention, it seems plausible that informational barriers would impose real constraints on sorting (this argument has been made in more detail by Merges (2001)).²⁷ In particular, the “patentability” of applications is difficult to assess *ex ante*, and there is no evidence that supervisory patent examiners attempt to do so before assigning applications to particular examiners.

5.2 Measure of examiner leniency

The institutional context outlined above motivates — conditional on Art Unit-by-application year fixed effects — the following candidate instrumental variable for application i filed in year t assigned to examiner j in Art Unit a :

$$Z_{itja} = \left(\frac{n_{\text{grant}_j} - \mathbf{1}(\text{grant}_i = 1)}{n_{\text{application}_j} - 1} \right)$$

²⁷Merges (2001) also argues that although sorting of patent applications to examiners may be efficient, an additional barrier to such sorting is the strong “all patents are created equal” tradition at the USPTO, which cuts strongly against any mechanism for separating and sorting patents.

where n_{grant_j} is the number of patents granted by examiner j , $n_{\text{application}_j}$ is the number of patent applications reviewed by examiner j , and grant_i is an indicator equal to 1 if patent application i was granted a patent. This expression measures the patent grant rate of examiner j , to which application i was assigned, for all gene patent application cases except application i itself. Motivated by the institutional context, we include Art Unit-by-application year fixed effects in our regressions, so that this instrument measures an examiner’s patent grant propensity relative to other examiners reviewing applications in that Art Unit in that application year. In our descriptive statistics and graphical analyses, we residualize the instrument by these Art Unit-by-application year fixed effects in order for the descriptive statistics and figures to reflect the same variation as our regression analyses.

For our baseline specification, we calculate this instrument based on the patent grant status of each gene patent application as of 2010. One concern is that this measure is censored for patent applications that are still in the examination process. To investigate the relevance of censoring, Appendix Figure D.2 plots — for our first four cohorts of gene patent applications, filed in 2001-2004 — the distribution of time between the application filing and the patent grant date, for the set of granted patents in our full sample. The average time to patent grant is around four years, and the cumulative share flattens out around six years. Given that we observe patent grant outcomes through 2010, this suggests that focusing on the subsample of gene patent applications filed in 2001-2004 should be relatively free of censoring concerns.

A separate concern is that if we only observed a small number of applications per examiner, our estimate of the variation in leniency across examiners would be overstated. To address this concern, we limit the sample to examiners that saw at least ten gene patent applications (Heckman, 1981; Greene, 2001). Together with two additional restrictions — Art Unit-years with at least two examiners, and Art Unit-years reviewing at least 10 applications — this restriction drops ~3% of our sample.

Appendix Figure D.3(a) presents a histogram of the residualized instrument Z_{itja} in the full sample. The standard deviation of the instrument is 0.10, with a range from -0.56 to 0.64. Appendix Figure D.3(b) presents histograms of the residualized instrument Z_{itja} in the human gene sample, as estimated in the full sample and (separately) the human gene sample. When estimated on the full sample, the standard deviation of the instrument is 0.09, with a range from -0.33 to 0.35; when estimated on the human gene sample, the standard deviation and range are similar.

5.3 Graphical analysis

We first present a simple graphical analysis that provides an analog to Figure 2. While our main analysis in this section will exploit the more continuous sources of variation available in our data, these graphs provide a more transparent way of displaying our main results graphically.

Figure 3 is constructed as follows. We start with the human gene sample of patent applications (N=1,582). Because patent applications can claim intellectual property over more than one human gene, and one human gene can be claimed in more than one patent application, we expand this data to patent application-human gene level observations. We classify each observation as being assigned to an above (triangle-denoted solid blue series) or below (circle-denoted dashed red series) median examiner based on the value of the examiner instrument for that application, as estimated in the full sample. To approximate the natural experiment described in Section 5.1, we condition on art unit-by-application year fixed effects. We then collapse this data into two groups, by above-median and below-median examiner leniency.

Figure 3(a) illustrates the time pattern of when each group receives patent grants over time. Consistent with the first stage estimates we document in Section 5.4, patent applications assigned to more lenient examiners are more likely to be granted patents. Consistent with the selection analysis we present in Section 5.5, comparing the pre-2001 levels and trends of both scientific publications (Figure 3(b)) and clinical trials (Figure 3(c)) suggests little evidence of selection: genes assigned to more lenient examiners appear to have similar levels of follow-on innovation to genes assigned to less lenient examiners prior to the application filing date. Finally, consistent with our estimates in Section 5.6, although these two groups of genes diverge in whether they are claimed by patents, we do not see any evidence of a divergence in follow-on innovation outcomes after these patents start to be granted (post-2001). That is, just as we saw in our comparison of successful and unsuccessful patent applications, these figures suggest that gene patents have not had a quantitatively important effect on either follow-on scientific research or on follow-on commercialization.

5.4 First stage estimates

Figure 4(a) presents a binned scatterplot relating our instrument (residualized by Art Unit-by-application year fixed effects) to the probability of patent grant in our full sample of patent applications (focus for now on the dark red plot; we will discuss the lighter yellow overlaid plot below). Visually, there is a strong relationship. Figure 4(b) plots this relationship in our human gene sample of patent applications, binning application-gene level data. The slope of the estimated linear fit lines in the two graphs are quite similar.

To quantify the magnitude of this upward slope, we estimate the following equation for a patent application i examined by patent examiner j filed in year t assigned to Art Unit a :

$$\mathbf{1}(\text{patent grant})_{itja} = \alpha + \beta \cdot Z_{itja} + \Sigma_{at} \mathbf{1}(\text{art unit})_{ta} + \varepsilon_{itja}$$

where the outcome variable $\mathbf{1}(\text{patent grant})_{itja}$ is an indicator variable equal to one if patent application i was granted a patent, Z_{itja} is our leave-one-out-mean (LOOM) gene patent grant rate instrument (as defined in Section

5.2), and $\Sigma_{at}\mathbf{1}(\text{art unit})_{ta}$ are a set of Art Unit-by-application year fixed effects. We estimate heteroskedasticity robust standard errors, clustered at the examiner level.

Column (1) of Panel A of Table 3 presents the estimate which corresponds to Figure 4(a). The point estimate implies that a 10 percentage point increase in an examiner's average patent grant rate is associated with a 6.8 percentage point increase in the likelihood that a patent application is granted a patent. The F-statistic is on the order of 400, well above the rule of thumb for weak instruments (Stock, Wright and Yogo, 2002). As a robustness check, Column (2) replaces our Art Unit-by-year fixed effects with Art Unit-by-year-by-class-by-subclass fixed effects, on the subsample for which these finer fixed effects can be estimated. The point estimate in Column (2) is very similar to and not statistically distinguishable from the point estimate in Column (1), suggesting that at least in our context, variation in the measured leniency of different examiners is unlikely to be generated by the systematic sorting of patent applications that are in more versus less 'patentable' technology classes or subclasses. Given the similarity of the estimates across Columns (1) and (2), we focus on Art Unit-by-year fixed effects for the remainder of the analysis.

In our sample of human gene patent applications — estimated on a dataset of unique application-level observations — we obtain a first stage coefficient of 0.865 (Column (1) of Panel B of Table 3) — slightly larger in magnitude although less precise than the corresponding estimate in Panel A. If we expand our human gene patent application sample to patent application-gene level observations — a change which gives more weight to patent applications which claim intellectual property over more genes — we obtain a first stage coefficient of 0.563 (Column (1) of Panel C of Table 3). In both Panel B and Panel C, the estimated coefficients are smaller in magnitude when we condition on the finer fixed effects in Column (2), although again the estimated coefficients are not statistically distinguishable from their counterparts in Column (1). As expected given the smaller sample of applications in the human gene samples, the estimates in Panel B and Panel C are less precise than the estimates in Panel A. The key first stage coefficient that will provide the basis for our empirical estimate is in Column (1) of Panel C; while statistically significant at the 1% level, the F-statistic is on the order of 8, just below the conventional cut-off of 10 for weak instruments. In ongoing work, we are expanding our sample of human gene patent applications to post-2005 years, which should increase the precision of this estimate.

As discussed in Section 5.2 our patent grant outcome is measured as of 2010 and is censored for patent applications that are still in the process of being examined, but this censoring should be less of a concern for our four early cohorts of gene patent applications (2001-2004). The point estimates on this sample are very similar to our baseline point estimates (results not shown), suggesting that censoring appears to not substantively affect the magnitude of the estimated first stage coefficient. Note that this is likely because the Art Unit-by-application year fixed effects largely account for differences in the probability of patent grant that are mechanically related to time since application. Given this similarity, we retain all cohorts of gene patent applications to retain a larger sample

size.

5.5 Investigating selection

In order for the examiner leniency measure constructed in Section 5.2 to be a valid instrumental variable for the likelihood that a given gene patent application is granted a patent, it must satisfy the exclusion restriction: the instrument can only affect follow-on innovation outcomes through the likelihood that a gene is patented. The institutional details described in Section 5.1 suggest that the assignment of applications to examiners is plausibly random — conditional on Art Unit-by-application year fixed effects — lending some a priori credibility to the exclusion restriction. In this section, we aim to empirically assess whether this assumption is reasonable by investigating whether patent applications assigned to ‘lenient’ and ‘strict’ examiners look similar on observable characteristics fixed at the time of patent application.

Ideally, we would empirically assess selection using variables that are correlated with the ‘patentability’ of the application at the time of filing, so that we could test whether applications that appear more patentable tend to be assigned to more lenient examiners. As discussed by Lemley and Sampat (2012), it is difficult to identify variables that measure the ‘patent-worthiness’ of an invention.²⁸ We here focus on two imperfect predictors of patentability: first, measures of the value of patent applications that are fixed at the time of application; and second, measures of the value of genes that are claimed in the patent application.

5.5.1 Selection tests based on the value of patent applications

A variety of metrics have been proposed as measures of the value of granted patents: forward citations (Trajtenberg, 1990), patent renewal behavior (Pakes, 1986; Schankerman and Pakes, 1986; Bessen, 2008), patent ownership reassignments (Serrano, 2010), patent litigation (Harhoff, Scherer and Vopel, 2003), and excess stock return values (Kogan, Papanikolaou, Seru and Stoffman, 2013). For our purposes, not all of these measures are appropriate: we need measures of patent value that are defined for patent applications (not just for granted patents), and also want a measure that is fixed at the time of application (and hence unaffected by subsequent grant decisions). For these reasons, we focus on two value measures which fit these criteria: patent family size and claims count.

Generally stated, a patent “family” is defined as a set of patent applications filed with different patenting authorities (e.g. US, Europe, Japan) that refer to the same invention. The key idea is that if there is a per-country cost of filing for a patent, firms will be more likely to file a patent application in multiple countries if they perceive the patent to have higher private value. Past work — starting with Putnam (1996) — has documented evidence that

²⁸In their paper, they show that two observable characteristics fixed at the time of application — the number of pages in the application and the patent family size — are not correlated with a measure of examiner experience (years of employment at the USPTO). That evidence provides some support for our exclusion restriction.

patent family size is correlated with other measures of patent value. We define patent family size as the number of unique countries in which the patent application was filed (as measured in the INPADOC patent family).

We use claims count as an alternative value measure that is fixed at the time of patent application, as proposed by Lanjouw and Schankerman (2001). The key idea here is that patents list “claims” over specific pieces of intellectual property, and that patents with more claims may be more valuable. Past work has documented mixed empirical evidence on whether that is a valid assumption based on correlations of claims counts with other value measures.

For our purposes, there are two key empirical questions we want to investigate using these measures. First, do patent family size and/or claims count predict patent grant? While clearly imperfect metrics of patentability, we document in Column 1 of Table 4 that these variables are predictive of patent grants for our full sample of patent applications (Panel A), for our human gene sample of patent applications (Panel B; here, family size is not statistically significant), and for our human gene patent application-gene level sample (Panel C). Second, is the predicted probability of patent grant — predicted as a function of family size and claims count — correlated with our examiner leniency instrument? We document in Column (2) of Table 4 that if we regress the predicted probability of patent grant (predicted as a function of family size and claims count) on our examiner leniency instrument (residualized by art unit-by-application year fixed effects), we estimate economically small (and — in two of three panels — ‘wrong signed’) statistically insignificant coefficients. This relationship for Panel A is displayed non-parametrically in the lighter yellow plot in Figure 4(a): consistent with the regression estimate, there is no visual relationship between the predicted probability of patent grant and our instrument.

5.5.2 Selection tests based on the value of claimed genes

We also investigate an alternative set of covariates: namely, characteristics of the genes that are included in the patent applications. Specifically, we use our measure of the level of scientific research activity on each gene in the years *prior* to the patent application (1970 to 1999) as a measure of the state of scientific knowledge about that gene.²⁹ Because one patent application often covers more than one human gene, this selection exercise relies on patent application-by-gene level data (since the covariates are defined at the gene level).

Do these gene-level covariates predict the ‘patentability’ of a given patent application? If we regress an indicator variable for whether a patent application is granted on these covariates in our sample of human gene patent applications, we find that these covariates are predictive of patent grants: the p-value from an F-test of the scientific

²⁹We also compile a host of other gene-level observable variables that measure a gene’s location on the genome — such as the chromosome on which a gene is located. Whereas the scientific research variables are defined for all observations in our sample, the genome location covariates are often missing; hence, we analyze these two sets of covariates separately so as to retain as many observations as possible, and focus on the scientific research variables.

research variables as predictors of patent grant is less than 0.001.³⁰ Hence, although these covariates are clearly imperfect proxies for patentability, they do appear to be useful for the purpose of this exercise. If we regress the predicted probability of patent grant (predicted as a function of our gene-level scientific research variables) on our examiner leniency instrument (residualized by art unit-by-application year fixed effects), we estimate an economically small and statistically insignificant coefficient of 0.0009 (s.e. 0.0053).³¹ This relationship is displayed non-parametrically in the lighter yellow plot in Figure 4(b): consistent with the regression estimate, there is no visual relationship between the predicted probability of patent grant and our instrument.

Taken together, the analysis in this section provides indirect support of our exclusion restriction in the following sense: we find no evidence that applications which appear more likely to be patented based on measures that are fixed at the time of filing are differentially sorted to more lenient examiners.

5.6 Reduced form and instrumental variables estimates

Table 5 presents our reduced form estimates, relating our examiner leniency instrument to follow-on innovation outcomes. Panel A of Table 5 pools all post-2001 years together; Panel B of Table 5 plots separate estimates by each individual year. For our scientific publication and clinical trial outcomes, we present OLS estimates with a log (plus one) outcome; for our diagnostic test outcome we present OLS estimates.

Interpreting the publications coefficient in Column 1 of Panel A of Table 5, a patent application assigned to a patent examiner with a 10 percentage point higher grant rate is associated with a statistically insignificant 0.3 percent decrease in the number of publications on genes claimed in that application. Our 95% confidence interval can reject declines of greater than 14.5%, and reject increases of more than 8.5%. Interpreting the clinical trials coefficient in Column 2 of Panel A of Table 5, a patent application assigned to a patent examiner with a 10 percentage point higher grant rate is associated with a statistically insignificant 0.4 percent decrease in the number of clinical trials on genes claimed in that application. Our 95% confidence interval can reject declines of greater than 14%, and reject increases of more than 5.6%. Our diagnostic test outcome in Column 3 of Panel A of Table 5 is marginally statistically significant ($p=0.072$), although the point estimate is quite small: a patent application assigned to a patent examiner with a 10 percentage point higher grant rate is associated with a .12 percentage point decrease in the number of clinical trials on genes claimed in that application. Our 95% confidence interval can reject declines of greater than .26 percentage points, and reject increases of more than .01 percentage points. Relative to a mean of 9%, this confidence interval suggests that we can reject declines in this outcome of greater than 3% and reject increases of greater than 0.1%.

³⁰This p-value is also less than 0.001 if we add the location covariates, and estimate this regression on the sub-sample of data with non-missing data on the location covariates.

³¹This estimate is not shown in Table 4, given that that table focuses on investigating selection based on patent application-level covariates.

Table 6 presents the analogous instrumental variables estimates, which simply scale our reduced form coefficients in Table 5 by the first stage estimate presented in Column 1 of Panel C of Table 3. The publications coefficient in Column 1 of Panel A of Table 6 implies that a patent application being granted a patent (instrumented by our examiner leniency instrument) is associated with a statistically insignificant 5.4 percent decrease in the number of publications on genes claimed in that application. Our 95% confidence interval can reject declines of greater than 25%, and reject increases of more than 14%. The 95% confidence intervals for our clinical trials outcome in Column 2 of Panel A of Table 6 can reject declines of greater than 25% and reject increases of more than 10%, and the 95% confidence interval for our diagnostic test outcome can reject declines of greater than 4.6 percentage points and reject increases of more than 0.2 percentage points. Relative to a mean of 9%, this confidence interval suggests that we can reject declines in this outcome of greater than 50% and reject increases of greater than 2.2%.

The estimates from our first quasi-experimental approach — comparing follow-on innovation across genes claimed in successful versus unsuccessful patent applications — are much more precise than these estimates, but have the disadvantage of relying on a less transparent source of variation. However, our confidence in interpreting those estimates as causal is strengthened by the fact that the examiner leniency instrument — which appears to generate a clean ‘experiment’ — generates similar results, albeit which are less precise. From an economic perspective, the estimates from our first approach can cleanly reject the effect sizes documented in the prior literature (Murray et al., 2008; Williams, 2013), whereas the estimates from our second approach can reject these effect sizes for some outcomes (publications and clinical trials) but not others (diagnostic tests).

5.7 Robustness

In this section, we document two additional robustness checks which aim to further assess the validity of our instrument.

5.7.1 Leave-one-out-mean (LOOM) initial gene patent grant rate instrument

Because patent grant decisions beyond the initial decision are in part a function of the applicant’s behavior, we examine data on initial decisions as a robustness check:

$$Z_{itja}^{\text{initial}} = \left(\frac{n_{\text{initial decision}=\text{grant}_j} - 1(\text{initial decision}_i = \text{grant})}{n_{\text{application}_j} - 1} \right).$$

Because most patent applications are initially issued a non-final rejection, these initial decision gene patent application grant rates are much lower than our Z_{itja} final grant rate measure.

This initial decision instrument is censored for patent applications that have not yet received an initial decision. In our sample about half of submitted applications receive an initial decision within two to three years, which

suggests that censoring is less of a concern when we focus on the early 2001-2004 gene patent application cohorts. Panel (a) of Appendix Figure D.4 is a scatterplot illustrating the correlation between our initial grant rate instrument $Z_{itja}^{\text{initial}}$ and final grant rate instrument Z_{itja} in our full sample; the two measures have a correlation of 0.47. We interpret the relationship between these variables as providing support for the assumption that our variation in grant rates represents a feature of the examiner rather than a feature of the applicant.

5.7.2 Non-gene patent grant rate instrument

As a second robustness check, we take advantage of the fact that most of the patent examiners who review gene patent applications in our sample also review non-gene patent applications. Because the non-gene patent application grant rate is not affected by a gene patent application, we can use a simple mean rather than a leave-one-out-mean:

$$Z_{itja}^{\text{non-gene applications}} = \frac{n_{\text{grant}_j}}{n_{\text{application}_j}}.$$

This variable measures the allowance rate of examiner j , to whom applicant i was assigned, for all non-gene patent applications. For this instrument, we again use the grant status of each gene application as of the end of our data (2010).

This alternative instrument is attractive for addressing the concern that — despite the qualitative background described above — gene patent applications could in theory be sorted to patent examiners in a way such that, e.g., less “innovative” gene patent applications were allocated to certain examiners, and those applications were both differentially likely to be granted and differed in our measures of follow-on innovation outcomes. To the extent that the Z_{itja} gene patent grant rate measure and $Z_{itja}^{\text{non-gene applications}}$ non-gene patent grant rate measure are strongly correlated, that would be consistent with the grant rate capturing a feature of the patent examiner, rather than simply reflecting something about the types of gene patent applications that were assigned to the examiner. Panel (b) of Appendix Figure D.4 is a scatterplot illustrating the correlation between Z_{itja} and $Z_{itja}^{\text{non-gene applications}}$. The two measures are strongly positively correlated: in the full sample the correlation is 0.68. We interpret the relationship between these variables as providing some support for the monotonicity assumption, given that examiners appear to make similar decisions when faced with different types of patent applications.

6 Discussion

While by necessity speculative, in this section we provide an interpretation of our empirical estimates in the context of the prior literature.³² Perhaps the most natural interpretation of the available empirical evidence to date is that intellectual property rights per se matter much less for follow-on innovation than do non-patent policies governing access to materials – such as provisions restricting “openness” in the sense of Murray, Aghion, Dewatripont, Kolev and Stern (2008), or imposing “practical excludability” in the sense of Cohen and Walsh (2008). This distinction is important because while theoretical models of intellectual property rights usually do not distinguish between legal versus practical excludability, the available empirical evidence suggests that this is a key policy-relevant distinction.

As in the model of Aghion, Dewatripont and Stein (2008), we can have in mind that basic research — largely by academics — generates research leads, which can then be commercialized by private firms. In their model, openness — defined as a researcher’s ability to exchange ideas or materials with other researchers — can foster basic research as well as the creation of new research lines. This idea that openness matters is supported by the empirical evidence in Murray et al. (2008) as well as the work of Furman and Stern (2011).

A key distinction that we wish to draw is the point emphasized by Cohen and Walsh (2008) in their distinction between “legal” and “practical” excludability: patents as generally used preserve open access to materials at the exploratory research stage (Walsh et al., 2003a,b). Academics report in surveys that they very rarely investigate whether research materials that they already have access to are covered by patents.³³ One interpretation of this is that “tolerated infringement” is a common feature of the use of patented materials by academic researchers (Walsh, Arora and Cohen, 2003b; Walsh, Cho and Cohen, 2005). In addition, from an enforcement perspective patent infringement may be hard to detect, and many private firms may openly tolerate infringement by academic researchers.³⁴

Once basic research produces new lines of research, private firms — if not the same firm as the patent holder — must negotiate licensing agreements, and hence in order to rationalize our empirical estimates it must also be the case that licensing markets operate relatively efficiently. Work by Walsh, Arora and Cohen (2003b) provides survey evidence consistent with relatively efficient licensing markets for biomedical patents: although the biomedical patenting landscape is complex, they document that few researchers report abandoning projects due to conflicts over patents.³⁵ Consistent with this combined explanation of tolerated infringement and relatively

³²Appendix C interprets our estimates in light of the recent US Supreme Court ruling on gene patents.

³³The Walsh, Cho and Cohen (2005) survey suggested that only about 5% of academic bench scientists reported regularly checking for patents on research inputs.

³⁴For example, in the famous case of Human Genome Sciences’ CCR5 patent, then-CEO William Haseltine publicly stated: “We would not block anyone in the academic world from using this for research purposes.”

³⁵Walsh, Cho and Cohen (2005) and Walsh, Cohen and Cho (2007) documented a similar finding in a larger survey.

efficient licensing markets, there has been very little gene patent-related litigation (Holman, 2008).

Two concrete examples may be helpful in providing more texture on this interpretation. First, in a now-classic paper Heller and Eisenberg (1998) documented that more than 100 issued US patents included the term 'adrenergic receptor' in the claim language, and pointed to this as an example of the complex biomedical patent landscape. However, in a comment on the Heller-Eisenberg paper published in the same issue of *Science*, Seide and MacLeod (1998) argued that based on a cursory patent clearance review, at most only a small number of licenses might be required. A similar point has been made in the context of gene patents by Holman (2012): although a large number of human genes are claimed by US patents, he argues that a reading of the claims suggests that few licenses would be required for common types of follow-on innovation. While anecdotal, these analyses suggest that documentation of biomedical patenting frequency should not necessarily be interpreted as evidence of a problem: analyses such as the one in this paper are needed in order to investigate whether these patents are actually affecting real economic behavior. Second, return to the example in the introduction, where pharmaceutical firm Bristol Myers reported abandoning research on more than fifty cancer-related proteins due to conflicts with gene patent holders (Pollack, 2001). This example was actually reported as part of a licensing agreement Bristol Myers brokered with another firm — Athersys — for a method enabling the use of protein targets without infringing gene patents. In the language of Walsh, Arora and Cohen (2003b), Athersys' technology would be an example of 'inventing around' such that patents may hinder follow-on innovation less than one might otherwise expect.

It is helpful to relate this interpretation to the results from three previous papers — Murray et al. (2008), Williams (2013), and Galasso and Schankerman (2013) — which documented evidence that intellectual property rights discourages follow-on innovation in some contexts. Murray et al. (2008) analyze a package of patents and other licensing restrictions that the private firm Dupont held on genetically engineered mice, and Williams (2013) analyzes a non-patent form of database protection that the private firm Celera held on their version of the sequenced genome. Importantly, both Dupont and Celera's forms of intellectual property restricted "openness": in order to gain access to Dupont's genetically engineered mice, or to Celera's sequenced genome, researchers needed to navigate legal contracts and restrictions. This feature stands in sharp contrast with how patents are commonly used in practice. For example, in our context, patented human genes were freely available to researchers with no restrictions — indeed, data on the entirety of the sequenced human genome was available throughout our time period in public open-access databases. In addition, in both the Dupont and Celera case studies, potential infringement was closely monitored (which was possible because researchers needed to contact Dupont and Celera directly in order to access the relevant research materials). Finally, as discussed in the introduction, our results are quite consistent with the recent contemporaneous contribution of Galasso and Schankerman (2013), which investigates how patents affect follow-on innovation using variation in patent invalidations generated by the quasi-random assignment of court cases to judges at the US Court of Appeals for the Federal Circuit. While they

document evidence that in some contexts patent invalidations increase citations to invalidated patents (suggesting that patents deter follow-on innovation) in their analysis of heterogeneous treatment effects across technology types, they find no evidence that patent invalidations increase follow-on innovation in the technology field which includes most of our gene patents.

Of course, this interpretation is at best suggestive, and given the available data it is not possible to rule out alternative explanations. For example, Lemley (2008) argues that both researchers and private firms simply ignore patents to a large degree, a view that is also consistent with our data. In order to test between these two theories, we would need to measure licensing agreements, data which is usually kept confidential.

7 Conclusion

The contribution of this paper is to investigate whether patents on human genes have affected follow-on scientific research and product development. Using administrative data on successful and unsuccessful patent applications submitted to the US Patent and Trademark Office (USPTO), we link the exact gene sequences claimed in each patent application with data measuring gene-related scientific research (publications) and commercial investments (clinical development). Building on this data, we develop two new sources of quasi-experimental variation: first, a simple comparison of follow-on innovation across genes claimed in successful versus unsuccessful patent applications; and second, use of the “leniency” of the assigned patent examiner as an instrumental variable for whether the patent application was granted a patent. Both approaches suggest that — on average — gene patents have not had quantitatively important effects on follow-on innovation.

This empirical evidence speaks against two existing views. First, there has been widespread concern from a policy perspective that patents on human genes may hinder follow-on innovation. For example, in the recent *Association for Molecular Pathology v. Myriad Genetics* case, the US Supreme Court invalidated patent claims on genomic DNA, arguing that such patents “would ‘tie up’ the use of such tools and thereby inhibit future innovation premised upon them.” Our empirical estimates do not provide support for patents hindering follow-on innovation in the context of human genes. Second, dating back at least to the academic work of Kitch (1977), many have argued that patents on basic discoveries play an important role in facilitating subsequent investment and commercialization. Again, our empirical estimates do not provide support for patents spurring follow-on innovation in the context of human genes.

Taken together with the prior literature, our evidence suggests two conclusions. First, for the case of human genes, the traditional patent trade-off of ex ante incentives versus deadweight loss may be sufficient to analyze optimal patent policy design, because the effects of patents on follow-on innovation appear to be quantitatively small. Second, the available evidence suggests there may be important heterogeneity in how and when intellectual

property rights have effects on follow-on innovation. A variety of evidence suggests that policies governing access to materials appear to have important effects on follow-on innovation (Murray et al., 2008; Furman and Stern, 2011; Williams, 2013), and the evidence from Galasso and Schankerman (2013) suggests that patents may deter follow-on innovation in some specific contexts. This growing body of empirical evidence suggests policy-relevant distinctions that should inform future theoretical and empirical research on cumulative innovation.

Of course, we are hesitant to generalize these conclusions to other markets given that in this paper we analyze the effects of patents on one specific type of technology (human genes). However, a broader contribution of this paper is to develop two new sources of quasi-experimental variation that may provide valuable opportunities to investigate the effects of patents in a variety of other applications. An analogy to the literature in labor economics investigating the disincentive effects of disability insurance on labor supply may be helpful. An influential early literature related labor supply outcomes to observational variation in disability insurance (Parsons, 1980), and subsequent papers provided refined approaches by drawing comparisons between accepted and rejected applicants (Bound, 1989; von Wachter et al., 2011) and using variation in the leniency of disability insurance examiners as an instrument for disability insurance receipt (Maestas, Mullen and Strand, 2013). Our paper provides an example of how each of these approaches can be applied to investigate the effects of patents, using administrative data on successful and unsuccessful patent applications and — in the case of the examiner leniency approach — information on the patent examiner assigned to review the patent application. Especially given the relatively small number of patent law changes in countries like the US in recent years, these two sources of quasi-experimental variation may provide valuable opportunities to investigate the effects of patents in a variety of other applications.

References

- Aghion, Philippe, Mathias Dewatripont, and Jeremy Stein**, “Academic freedom, private-sector focus, and the process of innovation,” *RAND Journal of Economics*, 2008, 39 (3), 617–635.
- Aizer, Anna, Shari Eli, Joseph Ferrie, and Adriana Lleras-Muney**, “The effects of childhood means-tested cash transfers on mortality,” 2013. unpublished mimeo, available <http://www.econ.ucla.edu/alleras/research/papers/childhood%20transfers%20v10-september%202013.pdf>.
- Anand, Bharat and Tarun Khanna**, “The structure of licensing contracts,” *Journal of Industrial Economics*, 2000, 48 (1), 103–135.
- Arora, Ashish**, “Licensing tacit knowledge: Intellectual property rights and the market for know-how,” *Economics of Innovation and New Technology*, 1995, 4 (1), 41–60.
- , **Andrea Fosfuri, and Alfonso Gambardella**, *Markets for Technology: The Economics of Innovation and Corporate Strategy*, MIT Press, 2001.
- Bacon, Neil, Doug Ashton, Richard Jefferson, and Marie Connett**, “Biological sequences named and claimed in US patents and patent applications: CAMBIA Patent Lens OS4 Initiative,” 2006. <http://www.patentlens.net> (last accessed 2 January 2012).
- Bessen, James**, “Holdup and licensing of cumulative innovations with private information,” *Economics Letters*, 2004, 82 (3), 321–326.
- , “The value of U.S. patents by owner and patent characteristics,” *Research Policy*, 2008, 37 (5), 932–945.
- Bound, John**, “The health and earnings of rejected disability insurance applicants,” *American Economic Review*, 1989, 79 (3), 482–503.
- Burk, Dan L.**, “Are human genes patentable,” *International Review of Intellectual Property and Competition Law*, 2013, 44 (7), 747–749.
- Caulfield, Timothy, Subhashini Chandrasekharan, Yann Joly, and Robert Cook-Deegan**, “Harm, hype and evidence: ELSI research and policy guidance,” *Genome Medicine*, 2013, 5 (21).
- Cockburn, Iain, Samuel Kortum, and Scott Stern**, “Are all patent examiners equal? Examiners, patent characteristics, and litigation outcomes,” in Wesley Cohen and Stephen Merrill, eds., *Patents in the Knowledge-Based Economy*, National Academies Press, 2003.
- Cohen, Wesley and John Walsh**, “Real impediments to academic biomedical research,” in Josh Lerner and Scott Stern, eds., *Innovation Policy and the Economy Volume 8*, University of Chicago Press, 2008.
- Doyle, Joseph J.**, “Child protection and child outcomes: Measuring the effects of foster care,” *American Economic Review*, 2007, 97 (5), 1583–1610.
- , “Child protection and adult crime: Using investigator assignment to estimate causal effects of foster care,” *Journal of Political Economy*, 2008, 116 (4), 746–770.
- Fehder, Daniel, Fiona Murray, and Scott Stern**, “Intellectual property rights and the evolution of scientific journals as knowledge platforms,” 2012. unpublished mimeo, available <http://www.hbs.edu/faculty/Lists/Events/Attachments/171/FMS%2001%2017%2012%20NBBER%20CONFERENCE%20FINAL.pdf>.
- Furman, Jeffrey and Scott Stern**, “Climbing atop the shoulders of giants: The impact of institutions on cumulative research,” *American Economic Review*, 2011, 101 (5), 1933–1963.

- Galasso, Alberto and Mark Schankerman**, “Patents and cumulative innovation: Causal evidence from the courts,” 2013. unpublished mimeo.
- Gans, Joshua, David Hsu, and Scott Stern**, “When does start-up innovation spur the gale of creative destruction?,” *RAND Journal of Economics*, 2002, 33 (4), 571–586.
- , —, and —, “The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays,” *Management Science*, 2008, 54 (5), 982–997.
- Golden, John M. and William M. Sage**, “Are Human Genes Patentable? The Supreme Court Says Yes and No,” *Health Affairs*, 2013, 32 (8), 1343–1345.
- Green, Jerry and Suzanne Scotchmer**, “On the division of profit in sequential innovation,” *RAND Journal of Economics*, 1995, 26 (1), 20–33.
- Greene, William**, “Estimating econometric models with fixed effects,” 2001. unpublished mimeo.
- Harhoff, Dietmar, F.M. Scherer, and Katrin Vopel**, “Citations, family size, opposition, and the value of patent rights,” *Research Policy*, 2003, 32 (8), 1343–1363.
- Harrison, Charlotte**, “Isolated DNA patent ban creates muddy waters for biomarkers and natural products,” *Nature Reviews Drug Discovery*, 2013, 12, 570.
- Heckman, James**, “The incidental parameter problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process,” in Charles Manski and Daniel McFadden, eds., *Structural Analysis of Discrete Data with Econometric Applications*, MIT Press, 1981.
- Heller, Michael and Rebecca Eisenberg**, “Can patents deter innovation? The anticommons in biomedical research,” *Science*, 1998, 280 (5364), 698–701.
- Holman, Christopher**, “The impact of human gene patents on innovation and access: A survey of human gene patent litigation,” *University of Missouri-Kansas City Law Review*, 2007, 76, 295–361.
- , “Trends in human gene patent litigation,” *Science*, 2008, 322 (5899), 198–199.
- , “Debunking the myth that whole-genome sequencing infringes thousands of gene patents,” *Nature Biotechnology*, 2012, 30 (3), 240–244.
- Jensen, Kyle and Fiona Murray**, “Intellectual property landscape of the human genome,” *Science*, 2005, 310 (5746), 239–240.
- Kesselheim, Aaron, Robert Cook-Deegan, David Winickoff, and Michelle Mello**, “Gene patenting - The Supreme Court finally speaks,” *New England Journal of Medicine*, 2013, 369 (9), 869–875.
- Kieff, Scott**, “Property rights and property rules for commercializing inventions,” *Minnesota Law Review*, 2001, 85, 697–754.
- Kitch, Edmund**, “The nature and function of the patent system,” *Journal of Law and Economics*, 1977, 20 (2), 265–290.
- Kling, Jeffrey**, “Incarceration length, employment, and earnings,” *American Economic Review*, 2006, 96 (3), 863–876.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman**, “Technological innovation, resource allocation, and growth,” 2013. MIT mimeo.

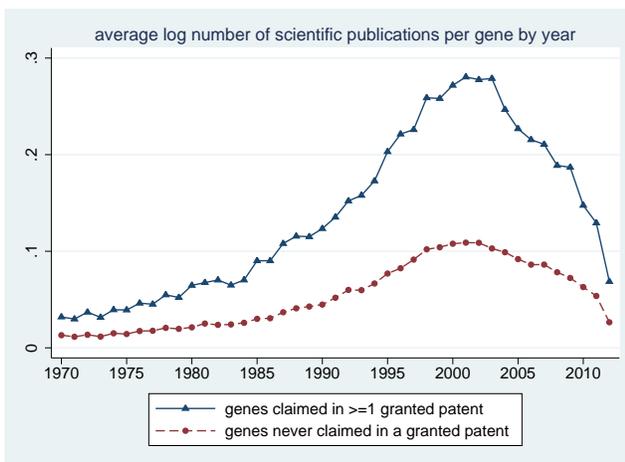
- Lanjouw, Jean and Mark Schankerman**, “Characteristics of patent litigation: A window on competition,” *RAND Journal of Economics*, 2001, 32 (1), 129–151.
- Lee, Byungwook, Taehyung Kim, Seon-Kyu Kim, Kwang H. Lee, and Doheon Lee**, “Patome: A database server for biological sequence annotation and analysis in issued patents and published patent applications,” *Nucleic Acids Research*, 2007, 35 (Database issue), D47–D50.
- Lemley, Mark**, “Rational ignorance at the patent office,” *Northwestern University Law Review*, 2001, 95 (4), 1495–1532.
- , “Ignoring Patents,” *Michigan State Law Review*, 2008, 19 (1), 19–34.
- **and Bhaven Sampat**, “Is the patent office a rubber stamp?,” *Emory Law Journal*, 2008, 58, 181–203.
- **and** —, “Examining patent examination,” *Stanford Technology Law Review*, 2010, 2.
- **and** —, “Examiner characteristics and patent office outcomes,” *Review of Economics and Statistics*, 2012, 94, 817–827.
- Lichtman, Doug**, “Rethinking prosecution history estoppel,” *University of Chicago Law Review*, 2004, 71 (1), 151–182.
- Maestas, Nicole, Kathleen Mullen, and Alexander Strand**, “Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt,” *American Economic Review*, 2013, 103 (5), 1797–1829.
- Marshall, Eliot**, “Lock up the genome, lock down research,” *Science*, 2013, 342, 72–73.
- Merges, Robert**, “Commercial success and patent standards: Economic perspectives on innovation,” *California Law Review*, 1988, 76, 803–876.
- , “As many as six impossible patents before breakfast: Property rights for business concepts and patent system reform,” *Berkeley Technology Law Journal*, 2001, 14, 577–615.
- Moon, Seongwuk**, “How does the management of research impact the disclosure of knowledge? Evidence from scientific publications and patenting behavior,” *Economics of Innovation and New Technology*, 2011, 20 (1), 1–32.
- Moser, Petra**, “Innovation without patents: Evidence from World’s Fairs,” *Journal of Law and Economics*, 2012, 55 (1), 43–74.
- Murray, Fiona and Scott Stern**, “Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis,” *Journal of Economic Behavior and Organization*, 2007, 63 (4), 648–687.
- , **Philippe Aghion, Mathias Dewatripont, Julian Kolev, and Scott Stern**, “Of mice and academics: Examining the effect of openness on innovation,” 2008. unpublished MIT mimeo.
- National Academy of Sciences**, *Reaping the Benefits of Genomic and Proteomic Research: Intellectual Property Rights, Innovation, and Public Health*, National Academies Press, 2006.
- Nordhaus, William**, *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*, MIT Press, 1969.
- Pakes, Ariel**, “Patents as options: Some estimates of the value of holding European patent stocks,” *Econometrica*, 1986, 54 (4), 755–784.

- Parsons, Donald**, “The decline in male labor force participation,” *Journal of Political Economy*, 1980, 88 (1), 117–134.
- Pollack, Andrew**, “Bristol-Myers and Athersys make deal on gene patents,” *New York Times*, 2001, 8 January.
- Putnam, Jonathan**, “The value of international patent protection,” 1996. Yale PhD dissertation.
- Rai, Arti and Robert Cook-Deegan**, “Moving beyond ‘isolated’ gene patents,” *Science*, 2013, 341, 137–138.
- Sampat, Bhaven**, “USPTO patent application data, Version 2,” 2011. <http://hdl.handle.net/1902.1/16402UNF:5:HFkba86yNYacbyMITt1KPVQ==BHAVENSAMPAT> (last accessed 2 January 2012).
- Schankerman, Mark and Ariel Pakes**, “Estimates of the value of patent rights in European countries during the post-1950 period,” *Economic Journal*, 1986, 96 (384), 1052–1076.
- Scherer, Frederic Michael**, “The economics of human gene patents,” *Academic Medicine*, 2002, 77 (12), 1348–1367.
- Scotchmer, Suzanne**, “Standing on the shoulders of giants: Cumulative research and the patent law,” *Journal of Economic Perspectives*, 1991, 5 (1), 29–41.
- , *Innovation and Incentives*, MIT Press, 2004.
- Seide, Rochelle and Janet MacLeod**, “Comment on Heller and Eisenberg,” 1998. ScienceOnline: <http://www.sciencemag.org/feature/data/980465/seide.dtl>.
- Serrano, Carlos**, “The dynamics of the transfer and renewal of patents,” *RAND Journal of Economics*, 2010, 41 (4), 686–708.
- Stock, James, Jonathan Wright, and Motohiro Yogo**, “A survey of weak instruments and weak identification in generalized method of moments,” *Journal of Business and Economic Statistics*, 2002, 20 (4), 518–529.
- Trajtenberg, Manuel**, *Economic Analysis of Product Innovation: The Case of CT Scanners*, Harvard University Press, 1990.
- United States Supreme Court**, “Association for Molecular Pathology et al. v. Myriad Genetics Inc. et al.,” 2013. 12-398.
- US General Accounting Office (GAO)**, “Intellectual property: USPTO has made progress in hiring examiners, but challenges to retention remain,” 2005.
- , “U.S. Patent and Trademark Office: Performance management processes,” 2010.
- USPTO**, “Utility examination guidelines,” *Federal Register*, 2001, 66 (4).
- von Wachter, Till, Jae Song, and Joyce Manchester**, “Trends in employment and earnings of allowed and rejected applicants to the Social Security Disability Insurance program,” *American Economic Review*, 2011, 101 (7), 3308–3329.
- Wade, Nicholas**, “A decade later, genetic map yields few new cures,” *New York Times*, 2010, 12 June.
- Walsh, John, Ashish Arora, and Wesley Cohen**, “Research tool patenting and licensing and biomedical innovation,” in Wesley Cohen and Stephen Merrill, eds., *Patents in the Knowledge-Based Economy*, National Academy Press, 2003.
- , —, and —, “Working through the patent problem,” *Science*, 2003, 299 (5609), 1021.

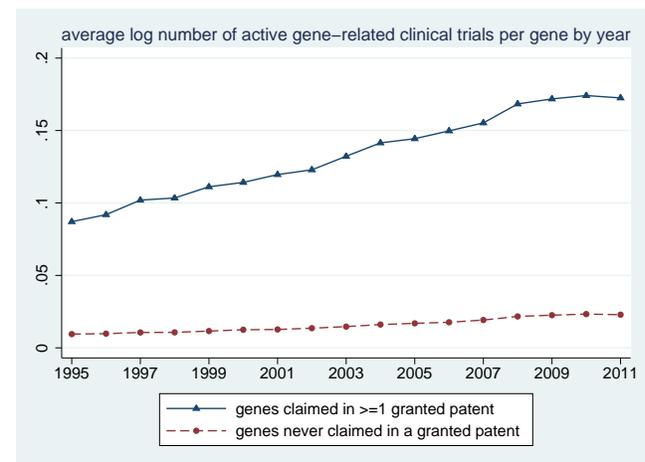
- , **Charlene Cho, and Wesley Cohen**, “View from the bench: Patents and material transfers,” *Science*, 2005, 309 (5743), 2002–2003.
- , **Wesley Cohen, and Charlene Cho**, “Where excludability matters: Material versus intellectual property in academic biomedical research,” *Research Policy*, 2007, 36 (8), 1184–1203.
- Williams, Heidi**, “Intellectual property rights and innovation: Evidence from the human genome,” *Journal of Political Economy*, 2013, 121 (1), 1–27.
- Wolinsky, Scott**, “An inside look at the patent examination process,” *The Metropolitan Corporate Counsel*, 2002, 10 (9), 18.

Figure 1: **Follow-on innovation on patented and non-patented human genes**

(a) Gene-level scientific publications

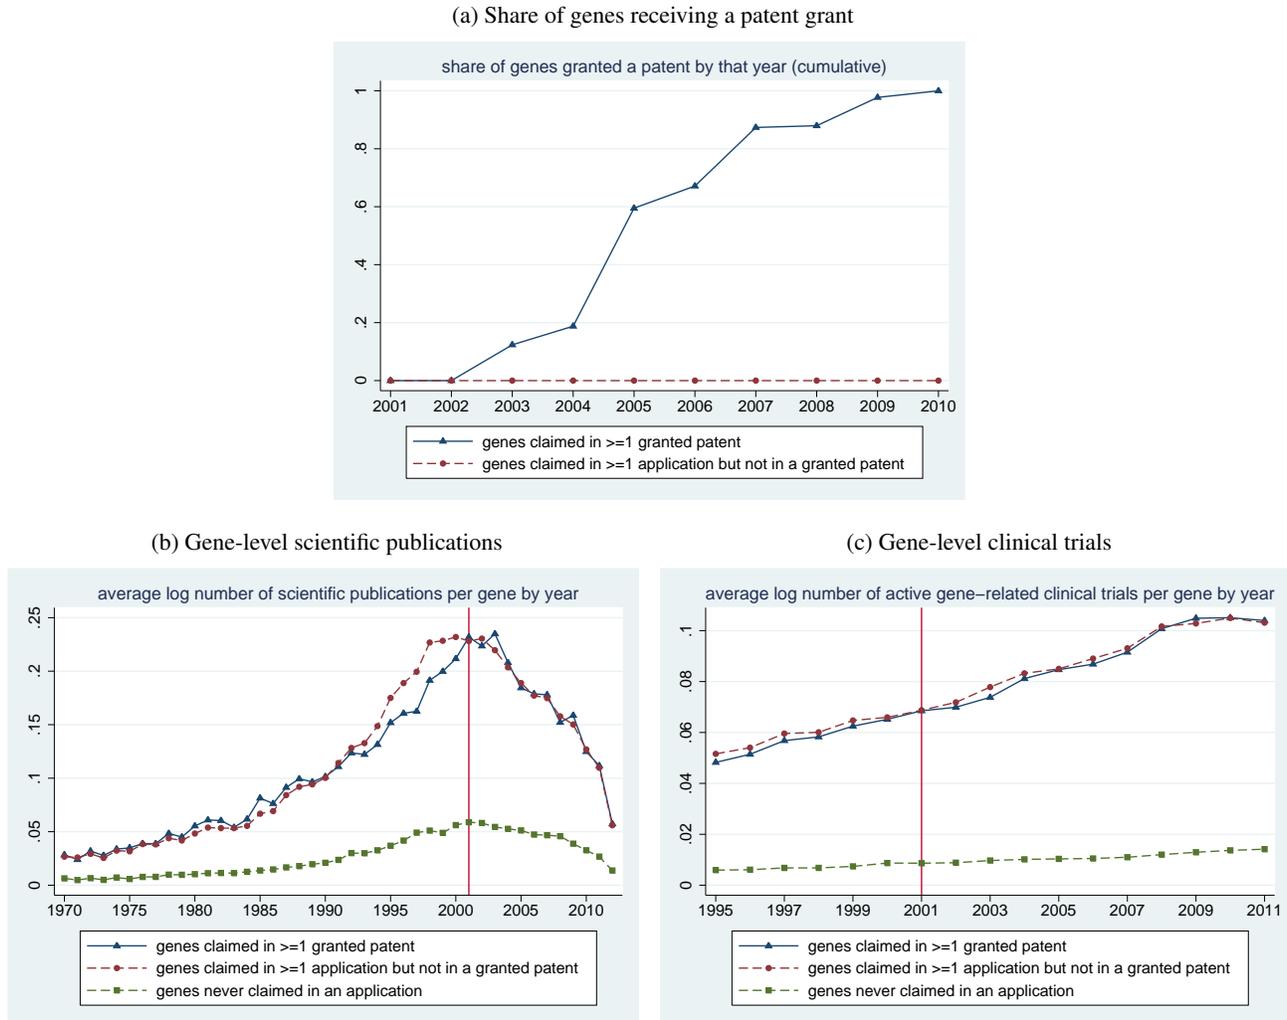


(b) Gene-level clinical trials



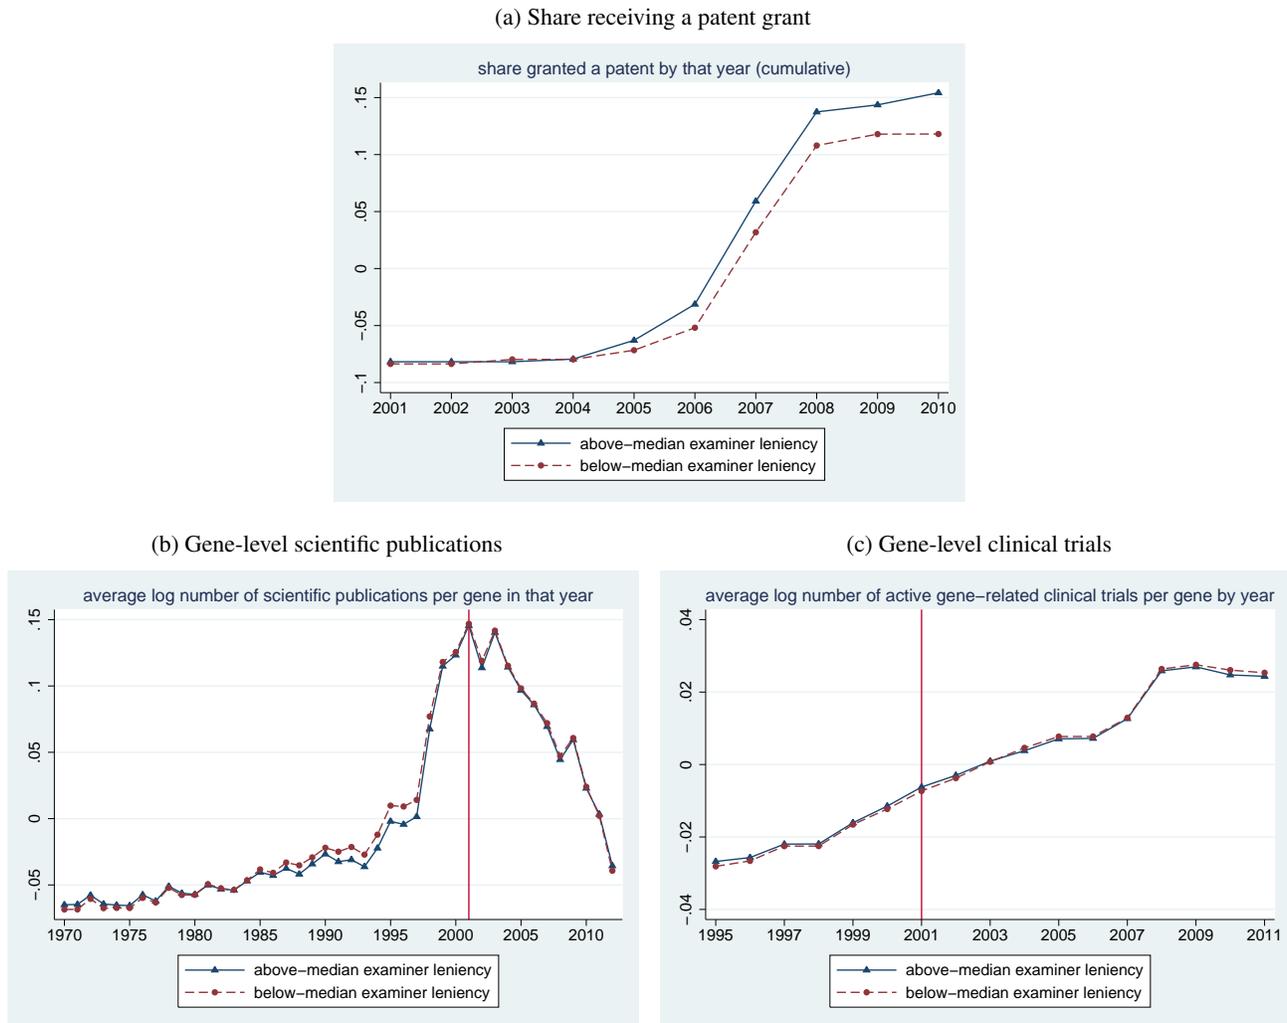
Notes: This figure plots trends in follow-on innovation by year separately for genes that ever receive a patent, and for genes that never receive a patent. The figure is constructed from gene-level data. Panel (a) uses gene-level scientific publications as a measure of follow-on innovation, and plots the average log number of scientific publications by year — plus one — in each year from 1970 to 2012. Panel (b) uses gene-level clinical trials as a measure of follow-on innovation, and plots the average log number of clinical trials by year — plus one — in each year from 1995 to 2011.

Figure 2: **Patents and follow-on innovation on human genes claimed in accepted/rejected patent applications**



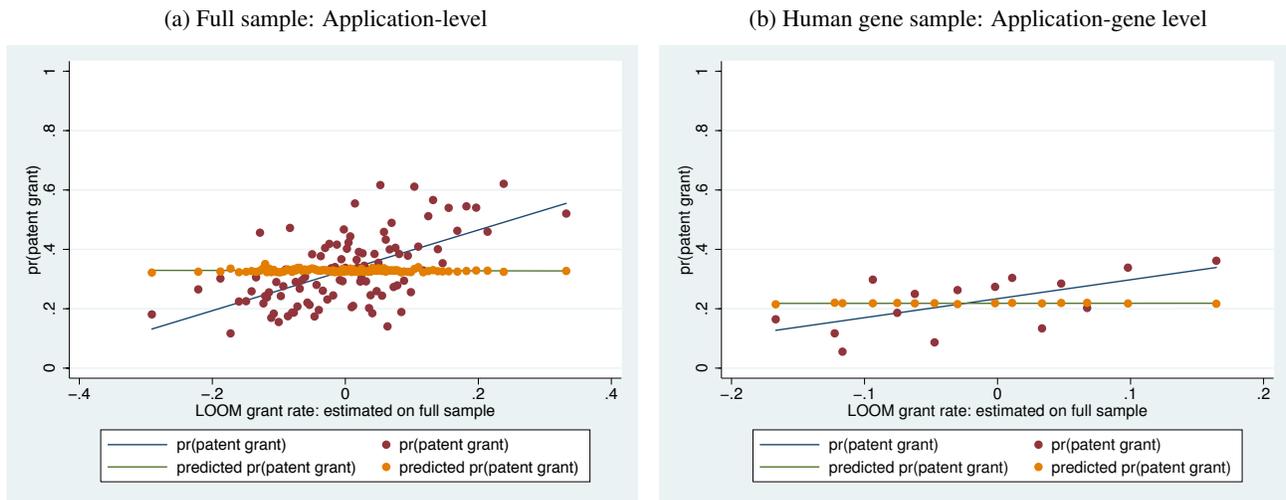
Notes: This figure plots trends in patenting and follow-on innovation by year separately for three groups of genes: genes claimed in at least one granted patent; genes claimed in at least one patent application but never in a granted patent; and (in Panels (b) and (c)) genes never claimed in a patent application. The figure is constructed from gene-level data. Panel (a) documents the share of genes receiving a patent grant by year; by construction, this is zero for the circle-denoted red dashed line in all years, is zero for the triangle-denoted blue line in 2001, and reaches one for the triangle-denoted blue line in 2010; the intermediate years simply illustrate the time path of patent grants between 2001 and 2010 for the triangle-denoted blue line. Panel (b) uses gene-level scientific publications as a measure of follow-on innovation, and plots the average log number of scientific publications by year — plus one — in each year from 1970 to 2012. Panel (c) uses gene-level clinical trials as a measure of follow-on innovation, and plots the average log number of clinical trials by year — plus one — in each year from 1995 to 2011. The vertical line in the calendar year 2001 in Panels (b) and (c) denotes that because this figure focuses on patents that were filed in or after 2001, all years prior to 2001 can be considered a pre-period and used to estimate the selection of genes into patenting based on pre-filing measures of scientific research (publications) and commercialization (clinical trials).

Figure 3: **Patents and follow-on innovation on human genes assigned to above/below leniency examiners**



Notes: This figure plots trends in patenting and follow-on innovation by year separately for two groups of genes: genes claimed in patent applications that were assigned to above-median leniency patent examiners, and genes claimed in patent applications that were assigned to below-lenient patent examiners. The figure is constructed from patent application-gene level data. Panel (a) documents the share of genes receiving a patent grant by year. Panel (b) uses gene-level scientific publications as a measure of follow-on innovation, and plots the average log number of scientific publications by year — plus one — in each year from 1970 to 2012. Panel (c) uses gene-level clinical trials as a measure of follow-on innovation, and plots the average log number of clinical trials by year — plus one — in each year from 1995 to 2011. The vertical line in the calendar year 2001 in Panels (b) and (c) denotes that because this figure focuses on patents that were filed in or after 2001, all years prior to 2001 can be considered a pre-period and used to estimate the selection of genes into patenting based on pre-filing measures of scientific research (publications) and commercialization (clinical trials).

Figure 4: **Probability of patent grant, by examiner leniency instrument (Z_{itja})**



Notes: This figure plots the first stage relationship between the probability of patent grant and our examiner leniency instrument, and plots estimates from two of our selection tests. Panel (a) is constructed from patent application-level data for our full sample of patent applications; Panel (b) is constructed from patent application-gene level data for our human gene sample of patent applications. Panel (a) presents a binned scatter plot relating our leave-one-out-mean (LOOM) grant rate instrument, residualized by Art Unit-by-application year fixed effects, to two variables: (1) the patent grant rate; and (2) the predicted patent grant rate, where we predict patent grant as a function of our two measures of patent value fixed at the time of application (patent family size and claims count). Panel (b) presents a binned scatter plot relating our LOOM grant rate instrument, residualized by Art Unit-by-application year fixed effects and estimated on the full sample, to two variables: (1) the patent grant rate; and (2) the predicted patent grant rate, where we predict patent grant as a function of our gene-level scientific research variables (scientific publications in each year from 1970 to 1999). The second series is similar if we also include gene location covariates as predictors, and estimate the prediction on the sub-sample of data with non-missing data on the location covariates.

Table 1: **Patent application-level summary statistics**

	<u>mean</u>	<u>std dev</u>	<u>min</u>	<u>max</u>	<u>obs</u>
<u>Panel A: full sample</u>					
application year	2004	2.61	1994	2010	30,243
0/1, patent granted as of 2010	0.33	0.47	0.00	1.00	30,243
<u>Panel B: human gene sample</u>					
application year	2002	0.72	1998	2005	1,582
0/1, patent granted as of 2010	0.25	0.43	0.00	1.00	1,582

Notes: This table shows summary statistics for our patent application-level data in each of our two samples: Panel A for the full sample of patent applications, and Panel B for the human gene sample of patent applications.

Table 2: **Follow-on innovation on human genes claimed in accepted/rejected patent applications**

<u>outcome variable:</u>	<u>log publications, 2001-12</u> (mean in levels: 4.84) (1)	<u>log clinical trials, 2001-11</u> (mean in levels: 0.25) (2)	<u>diagnostic test, 2012</u> (mean in levels: 0.12) (3)
<u>Panel A: Post-2001 outcomes</u>			
patent granted	-0.034 (0.018)*	-0.0036 (0.0156)	-0.0091 (0.0055)
<u>Panel B: Coefficients by year</u>			
patent granted*2001	0.0036 (0.0088)	-0.00022 (0.0063)	-
patent granted*2002	-0.0069 (0.0088)	-0.0019 (0.0064)	-
patent granted*2003	0.015 (0.0091)*	-0.0040 (0.0066)	-
patent granted*2004	0.0043 (0.0085)	-0.0020 (0.0069)	-
patent granted*2005	-0.0045 (0.0083)	-0.00031 (0.0071)	-
patent granted*2006	0.0017 (0.0079)	-0.0022 (0.0073)	-
patent granted*2007	0.0031 (0.0081)	-0.0015 (0.0074)	-
patent granted*2008	-0.0056 (0.0076)	-0.00087 (0.0079)	-
patent granted*2009	0.0085 (0.0077)	0.0020 (0.0080)	-
patent granted*2010	-0.0021 (0.0066)	0.00013 (0.0080)	-
patent granted*2011	0.0019 (0.0060)	0.00084 (0.0080)	-
patent granted*2012	0.0013 (0.0044)	-	-

Notes: This table estimates differences in follow-on innovation on genes claimed in at least one granted patent relative to genes claimed in at least one patent application but never in a granted patent. The sample for these regressions is constructed from gene-level data, and includes genes included in at least one patent application in our USPTO human gene patent application sample (N=15,530). Panel A documents differences in follow-on innovation outcomes pooled across all available post-2001 years: 2001-2012 for scientific publications (Column (1)); 2001-2011 for clinical trials (Column (2)); and 2012 for diagnostic tests (Column (3)). Panel B documents differences in these follow-on innovation outcomes for individual years (each row in Panel B presents an estimate from a separate year-specific regression). Estimates are from ordinary least squares models. Heteroskedasticity robust standard errors. *: p<0.10; **: p<0.05; ***: p<0.01.

Table 3: **First stage estimates: Patent grants by examiner leniency instrument**

<u>outcome variable:</u>	<u>0/1, =1 if patent granted</u>	
	(1)	(2)
<u>Panel A: full application-level sample ($\bar{y} = 0.33$)</u>		
LOOM grant rate	0.682 (0.033)***	0.675 (0.045)***
number of observations:	29,197	15,253
<u>Panel B: human gene application-level sample ($\bar{y} = 0.25$)</u>		
LOOM grant rate	0.865 (0.185)***	0.614 (0.248)**
number of observations:	1,566	1,148
<u>Panel C: human gene application-gene level sample ($\bar{y} = 0.22$)</u>		
LOOM grant rate	0.563 (0.200)***	0.472 (0.227)**
number of observations:	301,642	301,523
fixed effects:	art unit-year	art unit-year-class-subclass

Notes: This table presents first stage estimates, relating the probability of patent grant to our examiner leniency instrument. The sample for these regressions is constructed from patent application-level data in Panels A and B, and from patent application-gene level data in Panel C. Panel A documents estimates for our full sample of application-level data. Panel B documents estimates for our human gene sample of application-level data. Panel C documents estimates for our human gene sample of application-gene level data. Column (1) documents estimates that condition on Art Unit-by-application year fixed effects, and Column (2) replaces the Art Unit-by-application year fixed effects with Art Unit-by-application year-by-class-by-subclass fixed effects, estimated on the subsample of data meeting our sample restrictions. Heteroskedasticity robust standard errors are clustered at the examiner level. *: $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$.

Table 4: Predicted probability of patent grant by examiner leniency instrument

<u>outcome variable:</u>	<u>0/1, =1 if patent granted</u>	<u>predicted pr(patent grant)</u>
	(1)	(2)
<u>Panel A: full application-level sample</u> ($\bar{y} = 0.33$)		
family size	0.0027 (0.00079)***	
claims count	0.00026 (0.00013)*	
LOOM grant rate		-0.0031 (0.0037)
number of observations:	30,183	29,139
<u>Panel B: human gene application-level sample</u> ($\bar{y} = 0.25$)		
family size	0.0019 (0.0014)	
claims count	0.0018 (0.0005)**	
LOOM grant rate		-0.014 (0.033)
number of observations:	1,578	1,562
<u>Panel C: human gene application-gene level sample</u> ($\bar{y} = 0.22$)		
family size	0.0046 (0.0013)***	
claims count	0.0021 (0.00057)***	
LOOM grant rate		0.026 (0.042)
number of observations:	303,914	301,483

Notes: This table assesses whether patent application-level measures of value are predictive of patent grant, and whether the predicted probability of patent grant (predicted as a function of these value measures) is correlated with our examiner leniency instrument. The sample is constructed from patent application-level data in Panels A and B, and from patent application-gene level data in Panel C. Column (1) of this table regresses an indicator for the patent application being granted a patent on two proxy variables for patent value that are fixed at the time of application: family size and claims count. Column (2) predicts the probability of patent grant based on these two value measures, and regresses the predicted value on our examiner leniency instrument. Panel A documents estimates for our full sample of application-level data. Panel B documents estimates for our human gene sample of application-level data. Panel C documents estimates for our human gene sample of application-gene level data. The slightly larger sample sizes in Column (1) (relative to other tables) reflect that we do not impose our sample restrictions in this column, and the slightly smaller sample sizes in Column (2) (relative to other tables) reflect a small amount of missing data on the patent value measures. Estimates in Column (2) condition on Art Unit-by-application year fixed effects, estimated on the subsample of data meeting our sample restrictions. Heteroskedasticity robust standard errors are clustered at the examiner level. *: $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$.

Table 5: **Reduced form estimates: Follow-on innovation by examiner leniency instrument**

<u>outcome variable:</u>	<u>log publications, 2001-12</u> (mean in levels: 3.90) (1)	<u>log clinical trials, 2001-11</u> (mean in levels: 2.20) (2)	<u>diagnostic test, 2012</u> (mean in levels: 0.09) (3)
<u>Panel A: Post-2001 outcomes</u>			
LOOM grant rate	-0.030 (0.058)	-0.042 (0.049)	-0.012 (0.007)*
<u>Panel B: Coefficients by year</u>			
LOOM grant rate*2001	-0.016 (0.020)	-0.011 (0.013)	- -
LOOM grant rate*2002	-0.021 (0.024)	-0.011 (0.013)	- -
LOOM grant rate*2003	-0.015 (0.018)	-0.011 (0.015)	- -
LOOM grant rate*2004	-0.016 (0.018)	-0.015 (0.017)	- -
LOOM grant rate*2005	-0.022 (0.019)	-0.015 (0.018)	- -
LOOM grant rate*2006	-0.010 (0.019)	-0.014 (0.018)	- -
LOOM grant rate*2007	-0.016 (0.020)	-0.014 (0.018)	- -
LOOM grant rate*2008	-0.019 (0.019)	-0.015 (0.021)	- -
LOOM grant rate*2009	-0.024 (0.016)	-0.015 (0.020)	- -
LOOM grant rate*2010	-0.014 (0.018)	-0.017 (0.020)	- -
LOOM grant rate*2011	-0.007 (0.014)	-0.016 (0.022)	- -
LOOM grant rate*2012	-0.004 (0.012)	- -	- -

Notes: This table presents reduced form estimates, relating follow-on innovation outcomes to our examiner leniency instrument. The sample is constructed from patent application-gene level data, and includes patent application-gene level observations in our human gene sample (N=301,642). Panel A documents differences in follow-on innovation outcomes pooled across all available post-2001 years: 2001-2012 for scientific publications (Column (1)); 2001-2011 for clinical trials (Column (2)); and 2012 for diagnostic tests (Column (3)). Panel B documents differences in these follow-on innovation outcomes for individual years (each row in Panel B presents an estimate from a separate year-specific regression). Estimates are from ordinary least squares models and condition on Art Unit-by-application year fixed effects. Heteroskedasticity robust standard errors are clustered at the examiner level. *: p<0.10; **: p<0.05; ***: p<0.01.

Table 6: **Instrumental variables estimates: Follow-on innovation by examiner leniency instrument**

<u>outcome variable:</u>	<u>log publications, 2001-12</u> (mean in levels: 3.90)	<u>log clinical trials, 2001-11</u> (mean in levels: 2.20)	<u>diagnostic test, 2012</u> (mean in levels: 0.09)
	(1)	(2)	(3)
Panel A: Post-2001 outcomes			
patent granted (instrumented)	-0.054 (0.099)	-0.075 (0.088)	-0.022 (0.012)*
Panel B: Coefficients by year			
patent granted (instrumented)*2001	-0.029 (0.031)	-0.020 (0.023)	-
patent granted (instrumented)*2002	-0.036 (0.041)	-0.020 (0.023)	-
patent granted (instrumented)*2003	-0.026 (0.031)	-0.021 (0.025)	-
patent granted (instrumented)*2004	-0.029 (0.032)	-0.027 (0.029)	-
patent granted (instrumented)*2005	-0.038 (0.032)	-0.027 (0.030)	-
patent granted (instrumented)*2006	-0.019 (0.033)	-0.025 (0.032)	-
patent granted (instrumented)*2007	-0.029 (0.035)	-0.024 (0.032)	-
patent granted (instrumented)*2008	-0.034 (0.034)	-0.027 (0.037)	-
patent granted (instrumented)*2009	-0.042 (0.028)	-0.026 (0.035)	-
patent granted (instrumented)*2010	-0.025 (0.030)	-0.031 (0.036)	-
patent granted (instrumented)*2011	-0.012 (0.026)	-0.029 (0.038)	-
patent granted (instrumented)*2012	-0.006 (0.020)	-	-

Notes: This table presents instrumental variables estimates, relating follow-on innovation outcomes to patent grants instrumented with our examiner leniency instrument. The sample is constructed from patent application-gene level data, and includes patent application-gene level observations in our human gene sample (N=301,642). Panel A documents differences in follow-on innovation outcomes pooled across all available post-2001 years: 2001-2012 for scientific publications (Column (1)); 2001-2011 for clinical trials (Column (2)); and 2012 for diagnostic tests (Column (3)). Panel B documents differences in these follow-on innovation outcomes for individual years (each row in Panel B presents an estimate from a separate year-specific regression). Estimates are from instrumental variable models (instrumenting for patent grants with the examiner leniency instrument) and condition on Art Unit-by-application year fixed effects. Heteroskedasticity robust standard errors are clustered at the examiner level. *: p<0.10; **: p<0.05; ***: p<0.01.

A Appendix: Data construction

In this appendix, we describe our data construction in more detail. We first describe our patent application-related data, and then describe our human genome-related data.

Patent application-related data

Background: Application, publication, and patent numbers

A brief background on application, publication, and patent numbers is helpful for understanding our data on patent applications and granted patents from the United States Patent and Trademark Office (USPTO).

Application numbers: The USPTO assigns patent applications application numbers, which consist of a series code and a serial number.³⁶ The USPTO states that these application numbers are assigned by the Office of Patent Application Processing (OPAP) immediately after mail has been opened.³⁷ As suggested by this process, the USPTO and other sources note that application numbers are assigned chronologically.³⁸ While application serial numbers are six digits, the use and length of application series codes has changed over time: in recent years series codes are two digits, but previously these codes were one digit and historically series codes were not used.³⁹

Publication numbers: Traditionally, unsuccessful patent applications were not published by the USPTO. However, as part of the American Inventors Protection Act of 1999, the vast majority of patent applications filed in the US on or after 29 November 2000 are published eighteen months after the filing date. There are two exceptions. First, applications granted or abandoned before eighteen months do not appear in this sample unless the applicant chooses to ask for early publication. Lemley and Sampat (2008) estimate that about 17 percent of patents are granted before eighteen months, of which about half (46 percent) are published pre-patent grant. Second, applications pending more than eighteen months can “opt out” of publication if they do not have corresponding foreign applications, or if they have corresponding foreign applications but also have priority dates pre-dating the effective date of the law requiring publication (Lemley and Sampat, 2008).⁴⁰ If the patent application is published, then the USPTO assigns the application a publication number of the form USYEARXXXXXXX: a 2-digit country code, always US; followed by a 4-digit year (denoting year of publication); followed by a 7-digit identifier.

Patent numbers: Applications that are granted patents are assigned patent numbers. The number of characters in the patent number varies by the type of patent.⁴¹ Utility patent numbers are six or seven digits; reissue patents start with “RE” followed by six digits;⁴² plant patents start with “PP” followed by six digits; design patents start with “D” followed by seven digits; additions of improvement patents start with “AI” followed by six digits;⁴³ X-

³⁶For more details, see <http://www.uspto.gov/web/offices/ac/ido/oeip/taf/filingyr.htm>.

³⁷See <http://www.uspto.gov/web/offices/pac/mpep/s503.html>: “Application numbers consisting of a series code and a serial number are assigned by the Office of Patent Application Processing (OPAP) immediately after mail has been opened. If an application is filed using the Office’s electronic filing system, EFS-Web provides an Acknowledgement Receipt that contains a time and date stamp, an application number and a confirmation number.”

³⁸See <http://www.uspto.gov/web/offices/ac/ido/oeip/taf/filingyr.htm> (“In general, patent application serial numbers are assigned chronologically to patent applications filed at the U.S. Patent and Trademark Office.”) and http://www.thomsonfilehistories.com/docs/RESOURCES_Series_Codes.pdf (“US patent applications consist of a 2-digit series code and a 6-digit application serial that is assigned chronologically as they are received at the USPTO.”).

³⁹Note that design applications, provisional applications, and reexamination (*ex parte* and *inter partes*) applications are assigned different series codes; reissue patent application numbers follow the utility and design application structures. See <http://www.uspto.gov/web/offices/pac/mpep/s503.html> for details on these series codes.

⁴⁰For more details, see <http://www.uspto.gov/web/offices/pac/mpep/s1120.html> and the discussion in Lemley and Sampat (2010). Most applications not published eighteen months after filing are instead published sixty months after filing.

⁴¹For more details, see <http://www.uspto.gov/patents/process/file/efs/guidance/infopatnum.jsp>.

⁴²For more details on reissue patents, see <http://www.uspto.gov/web/offices/pac/mpep/s1401.html>.

⁴³Addition of improvement patents were issued between 1838 and 1861 and covered an inventor’s improvement on his or her own patented device. For more details, see §901.04 of <http://www.uspto.gov/web/offices/pac/mpep/s901.html>.

series patents start with “X” followed by seven digits;⁴⁴ H documents start with “H” followed by seven digits;⁴⁵ and T documents start with “T” followed by seven digits.⁴⁶

USPTO Patent Document Authority Files

The USPTO makes available Patent Document Authority Files.⁴⁷ The Pre-Grant and Patent Grant Authority Files contain listings of all US published applications beginning 15 March 2001 (“Pre-Grant”) and other published or granted patent documents (“Patent Grant”) found in USPTO Patent Image databases. Our versions of these files were downloaded on 24 March 2014 and are up to date as of February 2014.

Our version of the Pre-Grant Authority File includes 3,681,468 observations. Nearly all are utility patent applications (kind code=A1; 3,672,723; 99.76% of observations). Of the remaining 8,745 observations, 3,228 (kind code=A9; 0.09%) are corrections of published utility applications; 1,207 (kind code=A2; 0.03%) are subsequent publications of utility patent applications; and 4,310 (kind code=P1; 0.12%) are plant patent applications. One observation is denoted as the USPTO missing a copy of the relevant document images, and 1,025 observations are denoted as applications for which pre-grant publications were withdrawn. The identifier in the Pre-Grant Authority File is the publication number.

Our version of the Patent Grant Authority File includes 9,432,761 observations. The vast majority are utility patents (first character 0; 8,656,513; 91.77% of observations). Of the remaining 776,248 observations, 699,917 (first character D; 7.42%) are design patents; 44,742 (first two characters RE; 0.47%) are reissue patents; 24,257 (first two characters PP; 0.26%) are plant patents; 2,785 (first character X or first two characters RX; 0.03%) are X-series patents; 2,289 (first character H; 0.02%) are H documents; 1,968 (first character T; 0.02%) are T documents; and 290 (first two characters AI; <0.00%) are additions of improvement patents. Following the kind code in the US Patent and Trademark Office data files (listed as I2), we code re-issued X-series patents (which start with “RX” rather than “X”) as X-series patents.⁴⁸ There are 32,572 observations (0.35%) missing USPTO kind codes; 154 observations (<0.00%) denoted as the USPTO missing a copy of the relevant document images; and 32,572 observations (0.35%) denoted as applications for which pre-grant publications were withdrawn. The identifier in the Patent Grant Authority File is the patent number.

USPTO data on patent applications: *Cassis* BIB applications data

The USPTO’s *Cassis* BIB Applications DVD (Sampat, 2011) includes US patent applications published between 15 March 2001 and 31 December 2010. The data include 2,625,950 observations.

There are two identifiers in the *Cassis* data: publication numbers and application numbers. Publication numbers uniquely identify observations, whereas application numbers do not: 2,744 application numbers appear more than once in the dataset. Of those 2,744 application numbers, most (2,699) appear twice; 40 appear three times; and 5 appear four times. This is not unexpected, as it is known that application numbers can be published under multiple publication numbers.⁴⁹ Merging to the USPTO Pre-Grant Authority File reveals that almost all of these cases of non-unique application numbers are explained by subsequent publications of utility applications (697 observations) or by corrections of published utility applications (2,094 observations). The three remaining duplicate application numbers (8957425, 9726661, and 9728294) can be confirmed as having two application numbers on e.g. Google

⁴⁴X-series patents were issued between 1790 and 1836. For more details, see §901.04 of <http://www.uspto.gov/web/offices/pac/mpep/s901.html>.

⁴⁵H documents are part of the statutory invention registration series. For more details, see <http://www.uspto.gov/web/offices/ac/ido/oeip/taf/issudate.pdf>.

⁴⁶T documents are part of the defensive publication series. For more details, see <http://www.uspto.gov/web/offices/ac/ido/oeip/taf/issudate.pdf>.

⁴⁷Available at: <http://www.uspto.gov/patents/process/search/authority/index.jsp>.

⁴⁸For more details on kind codes, see <http://www.uspto.gov/patents/process/search/authority/kindcode.jsp>.

⁴⁹See, for example, WIPO Standard ST.1: <http://www.wipo.int/export/sites/www/standards/en/pdf/03-01-01.pdf>.

Patents, although the reason is not clear. We retain these duplicate application numbers in our sample in the absence of a clear reason to exclude them.

If we check the *Cassis* data against the USPTO Pre-Grant Authority File, as expected all observations in the USPTO's *Cassis* BIB Applications DVD appear in the Pre-Grant Authority File. As expected given that the *Cassis* BIB Applications DVD includes observations only through 31 December 2010 whereas the Pre-Grant Authority File includes observations through February 2014, many patent applications published between 1 January 2011 and February 2014 appear in the Pre-Grant Authority File but not in the *Cassis* BIB Applications DVD. In addition, 2,241 observations published prior to 1 January 2011 appear in the Pre-Grant Authority File but not in the *Cassis* BIB Applications DVD. These observations are roughly evenly distributed across publication years from 2001-2010; it is not clear why these observations are missing, but we are unable to include them in our analysis.

USPTO data on granted patents

In order to observe which US patent applications are granted patents, we use a USPTO administrative dataset covering patents issued from January 1, 1975 to December 31, 2010.⁵⁰ The data include 3,995,847 observations.

There are two identifiers in this data: application number and patent number. From the patent grant year variable in this data, we can infer whether and when a patent was granted as of December 31, 2010.

If we check these data against the USPTO Patent Grant Authority File, there are 12,770 reissue patents which appear in this USPTO grant data but not in the USPTO Patent Grant Authority File (none of the reissue patents in the USPTO grant data appear in the USPTO Patent Grant Authority File). There are 5,449,684 patents in the USPTO Patent Grant Authority File which do not appear in the USPTO grant data. Given that the USPTO Patent Grant Authority File includes all patent grants whereas the USPTO grant data includes only patents issued from 1975-2010, presumably most of these 5,449,684 patents were granted either prior to 1975 or after 2010, but as the USPTO Patent Grant Authority File does not include information on year of grant we have no direct way to verify that.

USPTO data on the patent application process: PAIR data

We also use data from the USPTO PAIR (Patent Application Information Retrieval) system. The PAIR data contains the key patent application-level variables needed for our analysis (*e.g.* examiner name). Google is currently crawling the universe of patent documents in the USPTO PAIR system and posts updated versions of the data as they become available; our version of this data is as of July 8, 2011.⁵¹ Not all of our patent applications were available in this Google data; we entered data on the remaining applications manually from the USPTO PAIR website.

Thomson Reuters Innovation database

We use the Thomson Innovation database as an additional source of data on patent applications. Specifically, the Thomson Innovation database provides information on two measures of patent value: claims count (as proposed by Lanjouw and Schankerman (2001)) and patent "family size" (as developed in Jonathan Putnam's 1996 dissertation). Below, we briefly describe these two value measures in more detail.

Claims count. We use claims count as a first proxy for the ex ante value of a patent application that is fixed at the time of patent application, as proposed by Lanjouw and Schankerman (2001). The key idea here is that patents list "claims" over specific pieces of intellectual property, so that patents with more claims may be more valuable. Past work has documented mixed empirical evidence on whether that is a valid assumption based on correlations of claims counts with other value measures.

⁵⁰Available at http://dvn.iq.harvard.edu/dvn/dv/boffindata/faces/study/StudyPage.xhtml;jsessionid=b8224569cfb3e22f8f305fcfdb51?globalId=hd1:1902.1/16412&studyListingIndex=0_b8224569cfb3e22f8f305fcfdb51; under the "Data & Analysis" tab, select "saved original (stata binary)" from the drop-down menu.

⁵¹Available at <http://www.google.com/googlebooks/uspto-patents-pair.html>.

Patent family size. We use patent “family” size as a second proxy for the ex ante value of a patent application that is fixed at the time of patent application, as proposed by Putnam (1996). A patent “family” is a group of related patents covering the same invention. Conceptually, this includes two types of patents: first, within-country families that may include continuations, continuations-in-part, and divisionals; and second, foreign filings covering the same technology in other jurisdictions. We here briefly describe each group of patents to motivate our family size measure:

- *Within-country patent families.* Within a country, patent families may include continuations, continuations-in-part, and divisionals. Because our focus is on US patent applications, we focus here on describing within-country patent families only for the US. This description summarizes material in the USPTO’s *Manual of Patent Examining Procedure*.⁵² A “continuation” is a subsequent application covering an invention that has already been claimed in a prior application (the “parent” application), before the parent application is abandoned or patented. The parent application is usually abandoned when the subsequent application is filed; continuation applications retain the original filing date and application number. A “continuation-in-part” is an application filed which repeats some portion of the parent application but also adds in new material not previously disclosed. Continuation-in-part applications are assigned a new application number, the original material retains the original filing date, and the new material receives a new filing date. A divisional application arises when an applicant divides claims in a parent application into separate patent applications. Each separate application is assigned a new application number, but retains the original filing date. Taken together, the use of continuations, continuations-in-part, and divisionals imply that more than one patent can issue from a single original patent application. Lemley and Sampat (2008) document that among utility patent applications filed in January 2001 and published by April 2006 (a sample of 9,960 applications), 2,016 “children” (continuations, continuations-in-part, or divisionals) had been filed by April 2006: around 30% were continuations, 20% were continuations-in-part, and 40% were divisionals (an additional 10% were of indeterminable types).
- *Foreign patent families.* Patent protection is jurisdiction-specific, in the sense that a patent grant in a particular jurisdiction provides the patent assignee with a right to exclude others from making, using, offering for sale, selling, or importing the invention into that jurisdiction during the life of the patent (subject to the payment of renewal fees). Hence, for any given patent application, applicants much choose how many jurisdictions to file patent applications in, and given that there is a per-jurisdiction cost of filing we would expect patents that are perceived by the applicant as more privately valuable to be filed in a larger number of jurisdictions.⁵³ The first patent application is referred to as the priority application, and the filing date of the the first application is referred to as the priority date; while the priority application can be filed in any jurisdiction, Putnam (1996) argues that the transaction costs involved with foreign filings (e.g. translation of the application) generally imply that domestic filing is cheaper than foreign filing, and that most priority applications are filed in the inventor’s home country. Under the Paris Convention for the Protection of Industrial Property (signed in 1883), all additional filings beyond the priority application that wish to claim priority to the priority application must occur within one year of the priority date. Putnam (1996) argues that most foreign applications — if filed — are filed near the one-year anniversary of the home country filing.
- *Commonly used measures of patent family size.* The term patent family can be used to describe different constructs: a patent family can be defined to include only documents sharing exactly the same priority or combination of priorities, or as all documents having at least one common priority, or as all documents that can be directly or indirectly linked via a priority document.⁵⁴ There are three commonly-used measures of family size which correspond to these three definitions: Espacenet (produced by the European Patent Office),

⁵²Available at <http://www.uspto.gov/web/offices/pac/mpep/s201.html>.

⁵³Multi-national routes such as applications filed with the European Patent Office or Patent Cooperation Treaty applications are intermediate steps towards filings in specific jurisdictions.

⁵⁴For details and examples, see <http://www.epo.org/searching/essentials/patent-families/definitions.html>.

the Derwent World Patents Index (DWPI; produced by Thomson Reuters), and INPADOC (produced by the European Patent Office).⁵⁵ Researchers tend to rely on these measures because collecting data from individual non-USPTO patent authorities would be quite time-consuming.

1. Espacenet uses a ‘simple’ patent family definition which defines a patent family as all documents with exactly the same priority or combination of priorities. This family is constructed based on data covering around 90 countries and patenting authorities.⁵⁶
 2. DWPI uses a similar patent family definition which defines a patent family as all documents with exactly the same priority or combination of priorities, but also includes non-convention equivalents (e.g. applications filed beyond the 12 months defined by the Paris Convention). This family is constructed based on data covering around 50 patent authorities and defensive publications (international technology disclosures and research disclosures). Continuations and divisionals are not included in the DWPI family definition.
 3. INPADOC defines a patent family more broadly, defining an ‘extended’ patent family as all documents that can be directly or indirectly linked via a priority document even if they lack a common priority. This family is constructed based on the same data as the Espacenet measure.
- *Our measure of patent family size.* For the purpose of our study, we would like to use the general concept of family size to develop a proxy for patent value that is fixed at the time the patent application is filed. Given this objective, it is clear that we should exclude continuations, continuations-in-part, and divisionals from our family size measure: these applications arise — by construction — after the filing date of the original patent application. In addition — and potentially more concerning in our specific context — the propensity for applications to develop continuations, continuations-in-part, or divisionals may differ across examiners, which could introduce an endogeneity problem. We define patent family size as the number of unique countries in which the patent application was filed (as measured in the INPADOC patent family).

DNA-related patent applications: CAMBIA Patent Lens data

The CAMBIA Lens database provides a list of published USPTO patent applications associated with human and non-human protein and DNA sequences appearing in patent claims (Bacon et al., 2006).⁵⁷ This data construction was supported by the Ministry of Foreign Affairs of Norway through the International Rice Research Institute for CAMBIA’s Patent Lens (the OS4 Initiative: Open Source, Open Science, Open Society, *Orzya sativa*).

Over the time period relevant for our analysis, US patent applications list DNA sequences in patent claims with a canonical ‘sequence listing’ label, e.g. SEQ ID NO:1 followed by the relevant DNA sequence. The CAMBIA Patent Lens data construction parses patent claims text for lists of SEQ ID NOs to determine which sequences are referenced in the claims. Importantly, CAMBIA makes available several versions of their data; following Jensen and Murray (2005), we focus on the dataset of nucleotide sequences (as opposed to amino acid sequences), and on the ‘in-claims’ subsample (as opposed to a larger dataset which includes DNA sequences listed in the text of the patent application, but not explicitly listed in the patent claims).

The CAMBIA Patent Lens data is updated over time; our version is current as of 8 February 2012. The level of observation is a patent-mRNA pair indexed by a publication/sequence number that combines the patent publication number and a mRNA sequence number. The patent publication numbers were extracted from the CAMBIA Patent Lens data using a Perl script.⁵⁸

⁵⁵For details on details on Espacenet, see <http://www.epo.org/searching/essentials/patent-families/espacenet.html>; for details on DWPI, see <http://www.epo.org/searching/essentials/patent-families/thomson.html>; and for details on INPADOC, see <http://www.epo.org/searching/essentials/patent-families/inpadoc.html>.

⁵⁶For a list, see [http://documents.epo.org/projects/babylon/eponet.nsf/0/2464E1CD907399E0C12572D50031B5DD/\\$File/global_patent_data_coverage_0711.pdf](http://documents.epo.org/projects/babylon/eponet.nsf/0/2464E1CD907399E0C12572D50031B5DD/$File/global_patent_data_coverage_0711.pdf).

⁵⁷Available at http://www.patentlens.net/sequence/US_A/nt-inClaims.fsa.gz.

⁵⁸We are very grateful to Mohan Ramanujan for help extracting this data from FASTA format. This Perl script is available on request.

Human gene-related patent applications: Patome data

Patome is a database server for biological sequence annotation and analysis in issued patents and published patent applications (Lee et al., 2007).⁵⁹ This data construction was supported by the Korean Ministry of Science and Technology (MOST).

Although the full Patome dataset contains issued patents and published patent applications from several jurisdictions — including Japan and Europe — in this paper we focus on the subsample of US published patent applications. As in the CAMBIA Patent Lens data, the Patome data construction parses patent application texts for lists of SEQ ID NOs to determine which sequences are referenced in patent applications. Following the methodology pioneered by Jensen and Murray (2005), BLAST (Basic Local Alignment Search Tool) searches are used to compare listed sequences against a census of potential matches in order to identify regions of similarity. Using these BLAST searches, the DNA sequences are annotated with mRNA and gene identifiers (RefSeq and Entrez Gene numbers).

The Patome data includes some patent applications which do not correspond to the definition of human gene patents proposed by Jensen and Murray (2005); to follow the Jensen-Murray definition, we impose some additional sample restrictions. First, the Patome data includes sequences which appear in the text of patent applications but are not explicitly listed in patent claims; to follow the Jensen and Murray (2005) definition of gene patents, we exclude observations which do not appear in the patent claims.⁶⁰ Second, following Jensen and Murray (2005) we limit the sample to BLAST matches with an E-value of exactly zero; the goal of this conservative E-value is to prevent spurious matches. Finally, following Jensen and Murray (2005) we limit the sample to disclosed sequences that are at least 150 nucleotides in length; the motivation of this restriction is that this is the average length of one human exon and yet still small enough to capture EST sequences.

As in Jensen and Murray (2005), many patents claim multiple variants of the same gene (that is, multiple mRNA sequences corresponding to the same gene). Following their methodology, we focus on variation in patenting across human genes.

Human genome-related data

Gene-level measures of scientific publications: OMIM data

We collect our measure of scientific research from the Online Mendelian Inheritance in Man (OMIM) database, a catalog of Mendelian traits and disorders. We use the full-text OMIM data and extract our variables of interest using a Python script.⁶¹ One gene can be included in more than one OMIM record, and one OMIM record can involve more than one gene. We tally the total number of publications related to each gene in each year across all OMIM records related to that gene.

Gene-level data on drug development: Pharmaprojects data

We collect data on gene-related drug development from the Pharmaprojects data.⁶² According to the company Citeline, which compiles and sells the Pharmaprojects data, “There is continual two-way communication between Pharmaprojects staff and their contacts in the pharmaceutical and biotechnology industries; both to gather new data and importantly, to verify information obtained from other sources.” Citeline employees gather drug information from company websites, reports, and press releases; annually every company covered in the database verifies information related to drugs in the development pipeline.

⁵⁹Available at http://verdi.kobic.re.kr/patome_int/.

⁶⁰Specifically, we merge the list of patent-mRNA numbers in the Patome data to the CAMBIA Patent Lens data, and drop observations that appear in Patome but not in the CAMBIA data; because our version of the CAMBIA data includes only patent-RNA observations listed in patent claims, this allows us to exclude Patome observations which are not explicitly listed in the claims of the patent applications.

⁶¹Available at <http://omim.org/downloads>.

⁶²Pharmaprojects data is available for purchase through Citeline or the Pharmaprojects website: <http://www.pharmaprojects.com>.

Pharmaprojects annotates a subset of the clinical trials in their data as related to specific Entrez Gene ID numbers. We construct a count of the number of clinical trials in which each gene is used as the basis for a pharmaceutical treatment compound in clinical trials in each year.

Gene-level data on diagnostic tests: GeneTests.org data

We collect our measure of genes being used in diagnostic tests from the GeneTests.org database.⁶³ This data includes a laboratory directory that is a self-reported, voluntary listing of US and international laboratories offering in-house molecular genetic testing, specialized cytogenetic testing, and biochemical testing for inherited disorders. US-based laboratories listed in GeneTests.org must be certified under the Clinical Laboratory Improvement Amendment of 1988, which requires laboratories to meet quality control and proficiency testing standards; there are no such requirements for non-US-based laboratories.

We use the GeneTests.org data as of September 18, 2012, which lists OMIM numbers for which there is any genetic test available in the Genetests.org directory. As with the OMIM data above, one gene can appear in more than one OMIM record, and one OMIM record can involve more than one gene. We construct an indicator for whether a given gene is used in any diagnostic test as of 2012.

Human genome-related covariates and crosswalks

We collect data on the two types of gene location covariates used by scientists: molecular location and cytogenetic location. Molecular location is recorded at the mRNA level; because a gene can link to more than one mRNA, we collapse this data to the gene level by taking an average of each variable (molecular location start, molecular location end, molecular location length, and orientation).⁶⁴ Cytogenetic location is recorded at the gene level, including map location variables (arm, region, band, and subband) and the chromosome.⁶⁵

NCBI-generated crosswalks are used to link OMIM numbers indexing genotype-phenotype pairs to Entrez Gene ID numbers,⁶⁶ and to match mRNA-level RefSeq accession/version numbers to Entrez Gene ID numbers.⁶⁷ We also use an NCBI-generated crosswalk that links discontinued Entrez Gene ID numbers to current Entrez Gene ID numbers, which is useful for linking Pharmaprojects observations that list discontinued Entrez Gene ID numbers to our data.⁶⁸

⁶³ Available at ftp://ftp.ncbi.nih.gov/pub/GeneTests/disease_OMIM.txt.

⁶⁴ Available at <ftp://ftp.ncbi.nih.gov/refseq/release/release-catalog/release54.accession2geneid.gz>.

⁶⁵ Available at ftp://ftp.ncbi.nih.gov/gene/DATA/GENE_INFO/Mammalia/.

⁶⁶ Available at ftp://ftp.ncbi.nih.gov/gene/DATA/mim2gene_partial.

⁶⁷ Available at <ftp://ftp.ncbi.nih.gov/refseq/release/release-catalog/release54.accession2geneid.gz>.

⁶⁸ Available at ftp://ftp.ncbi.nih.gov/gene/DATA/gene_history.gz.

B Appendix: The USPTO patent examination process

In this appendix, we describe the USPTO patent examination process in more detail.⁶⁹

B.1 Overview of the USPTO patent examination process

The USPTO is responsible for determining whether inventions claimed in patent applications qualify for patentability. The uniform mandate for patentability is that inventions are patent-eligible, novel, non-obvious, and useful.

Patent applications include a written description of the invention (the “specification”), declarations of what the application covers (“claims”), and a description of so-called prior art — ideas embodied in prior patents, prior publications, and other sources — that is known to the inventor and relevant to patentability. Once a patent application is received, as long as it satisfies a series of pre-examination formalities the Office of Patent Application Processing will assign it an application number, as well as a patent class and subclass.⁷⁰ These classifications, in turn, determine — based on a concordance between classes/subclasses and Art Units — the Art Unit to which the application is assigned, where Art Units are specialized groups of patent examiners that work on related subject matter.⁷¹ Once assigned to an Art Unit, a supervisory patent examiner (SPE) then refines the patent classification if it is incorrect. In some cases, this means the application needs to be re-assigned to another Art Unit, though that is thought to be rare.

The SPE then assigns the application to a patent examiner for review (via a process described in more detail below). While the patent application process up to this point is quite structured, from this point forward substantial discretion is left in the hands of individual examiners. In particular, the examiner is responsible for determining which — if any — claims in the application are patentable in light of prior art disclosed by the applicant as well as prior art found through the examiner’s own search. To give a sense of the time involved in the patent examination process, examiners have been estimated to spend an average of eighteen hours working on a given application (Lemley, 2001), although the process of evaluating a patent application takes several years.

As background for our empirical analysis, it is useful to outline the key steps in the examination process that occur from this point forward. The examiner sends a “first action” letter back to the applicant with an initial decision on the patent application. This initial decision can be an allowance, or (more commonly) a “non-final rejection.” For non-final rejections, applicants have a fixed length of time (usually six months) during which to revise the application. After receiving the applicant’s response, the examiner can then allow the application, negotiate minor changes, or send a second rejection. Counterintuitively, rejections in or after this second round are called “final rejections” but are in fact not final: there can be subsequent rounds of negotiation, and the only final resolution to a patent application that is not granted a patent is abandonment by the applicant (Lemley and Sampat, 2008).

B.2 Within-Art Unit assignment of patent applications to patent examiners

The process outlined above clarifies that the assignment of patent examiners to applications is a function of at least two factors: first, the Art Unit to which the application is assigned; and second, the year the application is filed, given that the group of examiners in an Art Unit will vary over time. In this section, we discuss how — within an Art Unit in a given application year — patent applications are assigned to patent examiners.

⁶⁹The discussion in this appendix draws heavily on Cockburn, Kortum and Stern (2003), Lemley and Sampat (2012), and GAO (2005), among other sources referenced in the text.

⁷⁰There are currently over 450 patent classes; the most common class in our sample is 435 (Chemistry: molecular biology and microbiology). There are currently more than 150,000 subclasses, which may be repeated among classes. For more details, see <http://www.uspto.gov/patents/resources/classification/overview.pdf>.

⁷¹There are over 300 Art Units; see <http://www.uspto.gov/patents/resources/classification/art/index.jsp>. For the current version of the class/subclass-to-Art Unit concordance, see <http://www.uspto.gov/patents/resources/classification/caau.pdf>. The main Art Units in our sample are from the 1600 group (Biotechnology and Organic Chemistry).

The USPTO does not publish rules regarding the assignment of applications within Art Units to particular examiners. Given this absence of formal written rules, Lemley and Sampat (2012) conducted written interviews with roughly two dozen current and former patent examiners and supervisory patent examiners to inquire about the assignment process. While the results of these interviews suggested that there is not a single “standard” assignment procedure that is uniformly applied in all Art Units, these interviews revealed no evidence of deliberate selection or assignment of applications to examiners on the basis of characteristics of applications other than those observed in standard USPTO datasets. In some Art Units, supervisors reported assigning applications to examiners based on the last digit of the application number; because application numbers are assigned sequentially in the central USPTO office, this assignment system — while not purposefully random — would be functionally equivalent to random assignment for the purposes of this study. In other Art Units, supervisors reported placing the applications on master docket based on patent classes and subclasses, with examiners specializing in those classes (or subclasses) being automatically assigned the oldest application from the relevant pool when requesting a new application. Consistent with what we would expect given these types of assignment mechanisms, Lemley and Sampat present empirical evidence that observable characteristics of patent applications are uncorrelated with characteristics such as the experience of the examiner to which the application was assigned. Unfortunately, we do not have information on the specific set of assignment processes used by the Art Units most common in our sample over the relevant time period.⁷² In the absence of such information, we rely on the interviews in Lemley and Sampat (2012) as a guide to designing our empirical specifications.

⁷²We have tried, unsuccessfully, to track down individuals who were supervisory patent examiners in the Art Units most common in our sample over the relevant time period.

C Appendix: Relevance of our estimates to the recent *AMP v. Myriad* case

As mentioned in the introduction, our estimates speak to the recent *AMP v. Myriad* case. The private firm Myriad Genetics was granted patent rights on human genes correlated with risks of breast and ovarian cancer. In 2009, the American Civil Liberties Union (ACLU) and the Public Patent Foundation filed suit against Myriad, arguing that many of Myriad's patent claims were invalid on the basis that DNA should not be patentable. One technical detail that is critical to understanding the *AMP v. Myriad* case is that two types of nucleotide sequences were at issue: naturally occurring genomic DNA (gDNA), and complementary or cDNA, which is produced in a laboratory using gDNA as a template. After a series of lower court decisions, in June 2013 the US Supreme Court issued a unanimous ruling drawing a distinction between these two types of sequences: "A naturally occurring DNA segment is a product of nature and not patent eligible...but cDNA is patent eligible because it is not naturally occurring."⁷³

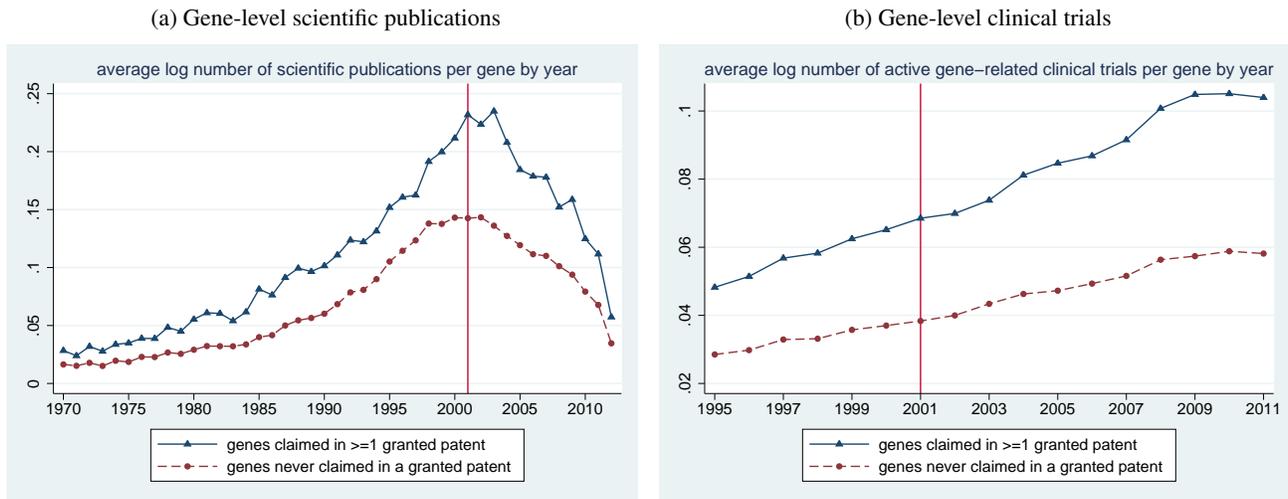
The question of whether DNA is patent eligible may at first blush seem very far removed from the economics of gene patents. Yet in fact, the US Supreme Court decision was made in part on the basis of the research question examined in this paper: whether patents on human genes would impede follow-on innovation. A brief background on patent eligibility is helpful in clarifying this point. The patent eligibility criteria set out in the US Constitution (35 U.S.C. §101) has long been interpreted to exclude laws of nature, natural phenomena, and abstract ideas from patent eligibility. The *AMP v. Myriad* decision followed this precedent, arguing that "[g]roundbreaking, innovative, or even brilliant" discoveries of natural phenomena should be patent-ineligible, because patents "would 'tie up' the use of such tools and thereby inhibit future innovation premised upon them." As discussed by Rai and Cook-Deegan (2013), the Court decision essentially aimed to draw a line between patent-eligible and patent-ineligible discoveries based on the "delicate balance" between patents prospectively creating incentives for innovation and patent claims blocking follow-on innovation. In the end, the Court drew this line by ruling naturally occurring DNA patent-ineligible, and non-naturally occurring cDNA patent-eligible.

Numerous legal scholars have argued that the distinction between DNA and cDNA is "puzzling and contradictory" (Burk, 2013) given that "both isolated sequences and cDNA...have identical informational content for purposes of protein coding" (Golden and Sage, 2013); in interviews, patent attorneys expressed similar confusion (Harrison, 2013). A recent analysis of gene patent claims by Holman (2012) concluded that most human gene patents claimed cDNA, and would thus be unaffected by the Court ruling.

⁷³The earlier decisions were a 2010 ruling by the US District Court for the Southern District of New York (see <http://www.pubpat.org/assets/files/brca/brcasjgranted.pdf>) and a 2011 ruling by the US Court of Appeals for the Federal Circuit (see <https://www.aclu.org/files/assets/10-1406.pdf>); a subsequent re-hearing of the case by the US Court of Appeals at the request of the US Supreme Court did not substantively change this decision.

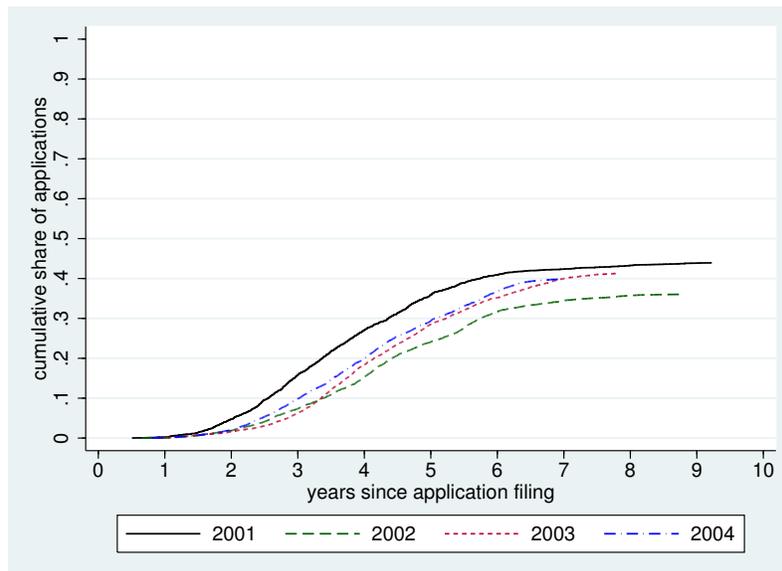
D Appendix: Additional results

Figure D.1: **Follow-on innovation on patented and non-patented human genes [patents filed in or after 2001]**

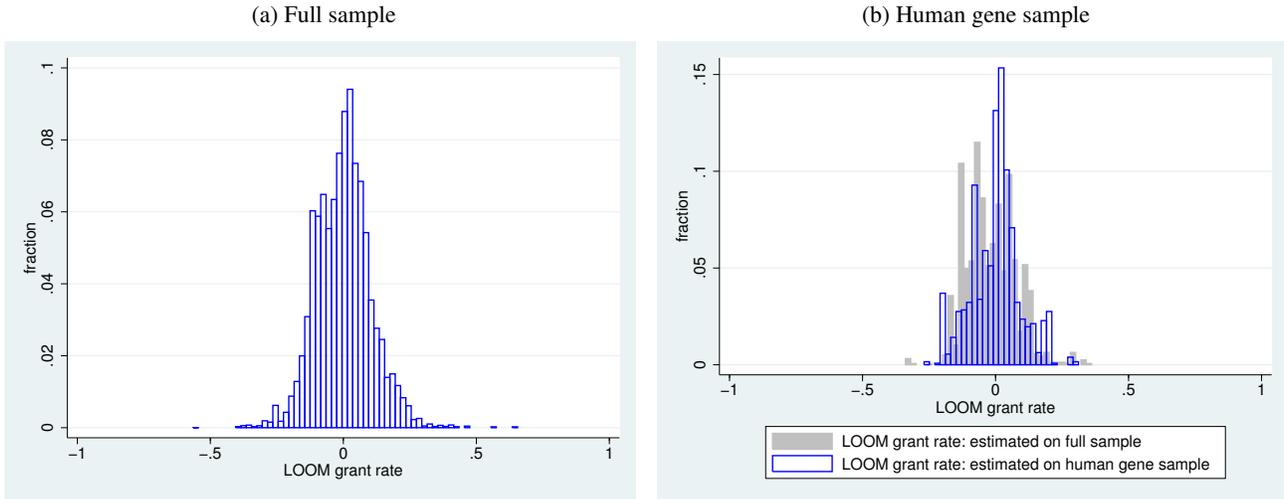


Notes: This figure plots trends in follow-on innovation by year separately for genes that ever receive a patent that was filed in or after 2001, and for genes that never receive a patent that was filed in or after 2001. Panel (a) uses gene-level scientific publications as a measure of follow-on innovation, and plots the average log number of scientific publications by year — plus one — in each year from 1970 to 2012. Panel (b) uses gene-level clinical trials as a measure of follow-on innovation, and plots the average log number of clinical trials by year — plus one — in each year from 1995 to 2011. The vertical line in the calendar year 2001 denotes that because this figure focuses on patents that were filed in or after 2001, all years prior to 2001 can be considered a pre-period and used to estimate the selection of genes into patenting based on pre-filing measures of scientific research (publications) and commercialization (clinical trials).

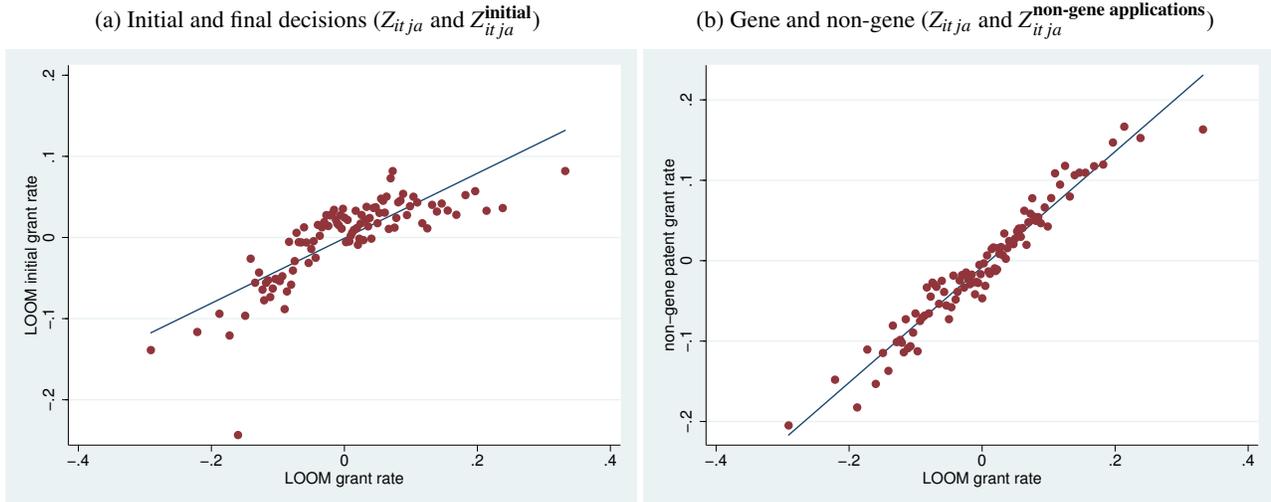
Figure D.2: **Cumulative share of applications granted a patent, by years since filing**



Notes: This figure plots — for our first four cohorts of patent applications, filed in 2001-2004 — the distribution of time between the application filing and the patent grant date, for the set of granted patents in our full sample.

Figure D.3: **Histogram of examiner leniency instrument (Z_{itja})**

Notes: This figure plots a histogram of our patent application-level observations by our leave-one-out-mean (LOOM) grant rate instrument, residualized by Art Unit-by-application year fixed effects. Panel (a) plots this variable for the set of patent applications in our full sample. Panel (b) plots two versions of this variable — one estimated on the full sample, and one estimated on the human gene sample — for the set of patent applications in our human gene sample.

Figure D.4: **Correlation between alternative examiner leniency instruments**

Notes: Panel (a) is a binned scatterplot relating our leave-one-out-mean (LOOM) grant rate instrument, residualized by Art Unit-by-application year fixed effects, to our LOOM initial grant rate instrument, residualized by Art Unit-by-application year fixed effects. Panel (b) is a binned scatterplot relating our LOOM grant rate instrument to our non-gene patent grant rate instrument, residualized by Art Unit-by-application year fixed effects. Both panels use the full sample set of patent applications.